



Optimizing Manufacturing with AI – Root Cause Analysis & Actionable Insights

This technical whitepaper showcases a real-world deployment of generative AI for Schaeffler using Microsoft Azure OpenAI Services.



Optimizing Manufacturing with AI

Root Cause Analysis & Actionable Insights

Celebal Technologies has developed an advanced Root Cause Analysis and Insights system, harnessing the power of generative AI for Schaeffler, a leading manufacturing company. The solution utilizes Azure AI Services and Azure OpenAI models to aggregate data from multiple sources, enabling precise identification of root causes and the recommendation of effective solutions. Additionally, it integrates with production APIs to retrieve real-time data and generate insightful summaries. With LLMOps integration and a user feedback mechanism, the solution continuously enhances response accuracy and reliability, driving smarter decision-making in manufacturing.

Understanding the Operational Disconnects

In large-scale manufacturing environments like Schaeffler, machines across multiple production lines, work centers, and workstations operate continuously. But when issues arise, identifying the root cause and finding a quick solution becomes a challenge, leading to inefficiencies and delays.

Key Challenges Across Manufacturing Workflows



Operator Flow

- Operators depend on manual communication to resolve issues
- Critical machine data is scattered across PCS, PMO, and Knowledge Graph, delaying diagnosis
- A real-time, data-driven system is needed for quick and accurate solutions.



User Manual Flow

- Machine setup requires digging through lengthy manuals, slowing troubleshooting
- A query-based system can surface relevant instructions instantly, saving time and effort.



Team Lead Flow

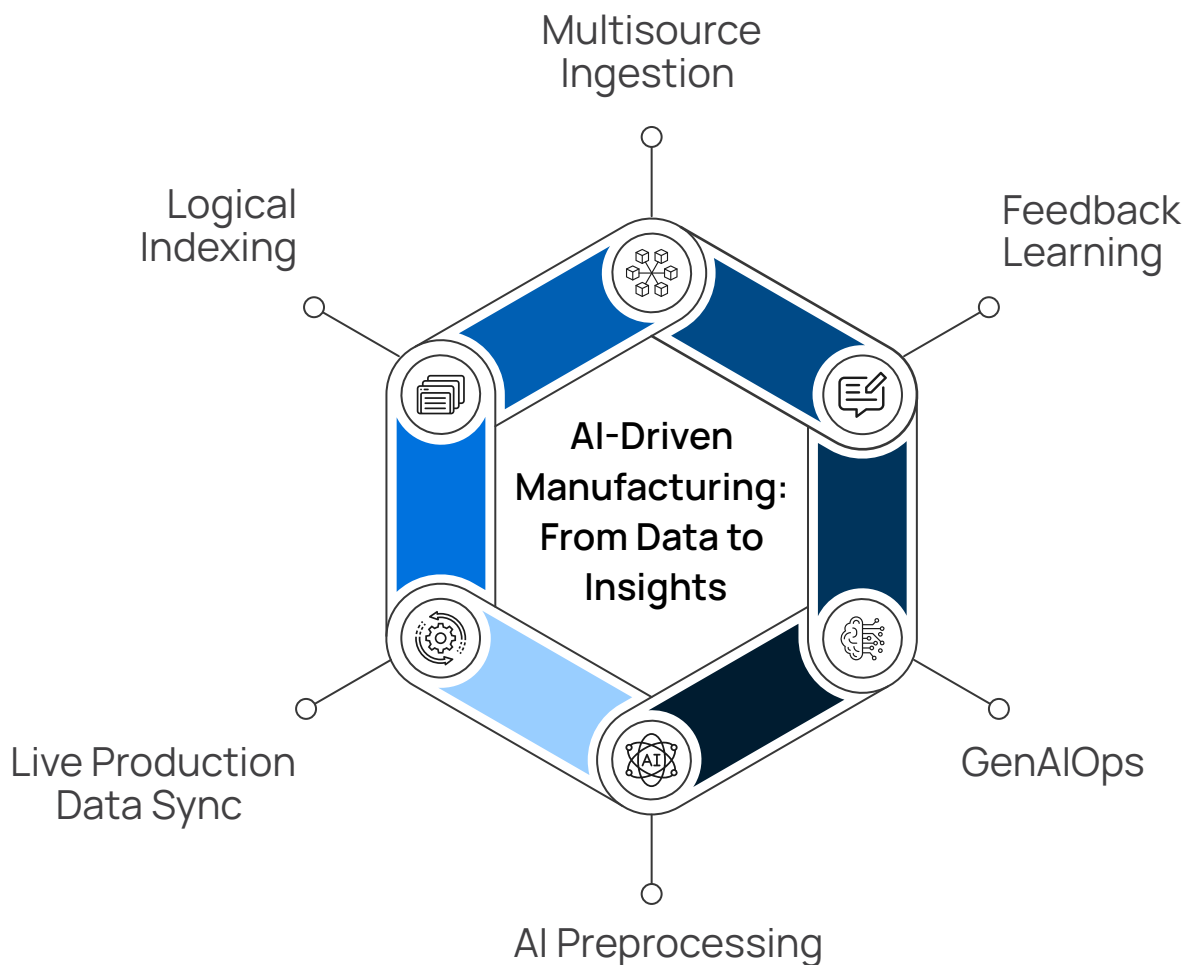
- Benchmarking needs data from multiple APIs, with manual field input in the UI
- Automating this process would speed up decision-making and improve optimization across operations.



Common Terms Flow

- Technical jargon causes confusion and inconsistency
- Currently, definitions are found manually
- An AI-powered glossary would bring clarity and faster issue resolution.

Revolutionizing Manufacturing Intelligence with Azure's AI-Powered Ecosystem



Manufacturing is a world of precision, where every second counts. When a machine breaks down or a production line slows, identifying the root cause quickly is the difference between a minor hiccup and a major disruption. That's where our AI-driven Root Cause Analysis and Insights Generator steps in, powered by Azure's cutting-edge AI services and Microsoft's intelligent search capabilities. This solution seamlessly integrates real-time data processing, generative AI, and automated insights to help Schaeffler tackle complex manufacturing challenges with speed and accuracy.

At the core of this system is a robust data ingestion pipeline that brings scattered manufacturing data into a centralized AI-powered search. Machine logs, production data, and maintenance records from PCS and PMO tables are preprocessed, with Azure OpenAI handling translations and error corrections. The processed data is indexed in Azure AI Search, ensuring instant access to critical insights.

To take things a step further, Microsoft's GraphRAG technology helps refine search results by linking relevant data points in a knowledge graph, making troubleshooting more accurate and intuitive. For machine manuals, Azure Document Intelligence extracts structured content, breaking it into larger ten-page chunks for context and smaller two-page chunks for precise retrieval. This means operators no longer have to waste time flipping through thick manuals they get exactly what they need, when they need it. Similarly, an AI-powered search index is created for manufacturing terminology, allowing users to quickly find definitions and explanations without confusion.

But collecting data is just one piece of the puzzle. The real magic happens in the bot pipeline, where AI transforms raw queries into actionable insights. Every question whether from an operator diagnosing a machine failure or a team lead analyzing production efficiency goes through multiple processing layers. Typos are corrected, translations applied, follow-up handled, guardrails enforced, and intent classified using Azure OpenAI's advanced models.

If an operator needs to pinpoint the cause of a breakdown, the system pulls relevant solutions using hybrid search across AI search indexes and GraphRAG, ensuring the most relevant information surfaces. For team leads monitoring production performance, the bot dynamically triggers the right APIs to retrieve benchmarking data, and then extracts key insights using prompt-based techniques. In the case of user manual queries, the bot smartly prioritizes retrieval using the big chunk index first to get a broad context, then refines results using the small chunk index for pinpoint accuracy. If a user asks about a technical term, the AI instantly retrieves the definition, ensuring clarity without delays.

Continuous Learning and Real Time Intelligence for Smarter Manufacturing

To keep operations at peak efficiency, the system continuously evaluates itself. Leveraging RAGAS scoring, it measures response accuracy, contextual relevance, and summarization quality. With every interaction, the system learns and improves refining its answers based on real-world usage and feedback. Built on Azure AI Services, OpenAI models, and Microsoft's advanced retrieval techniques, this solution is transforming how Schaeffler approaches manufacturing intelligence. No more guesswork. No more searching endlessly through documentation. Just fast, precise, and actionable insights that empower smarter decisions and smoother operations.

We've also implemented response caching, a significant breakthrough for performance and cost-efficiency. Every new query is checked against Azure Cosmos DB; if a similar one has already been processed, the system returns a pre-generated summary drastically reducing latency and server load.

To put users in control, we've added real-time feedback integration. A simple thumbs-up or thumbs-down lets the system know how it performed, triggering deeper analysis for ongoing improvement. Additionally, real-time streaming of responses makes interactions feel instant users see answers being generated live, enhancing both speed and engagement. Every update makes the system sharper, faster, and more intuitive.

This is AI in manufacturing adaptive, intelligent, and built for impact.

Redefining Speed and Precision in AI

We've pushed the boundaries of AI performance, achieving unparalleled accuracy and blazing-fast response times. Our typo-ambiguity handling now boasts an impressive 98.7% accuracy, while translation reaches a near-perfect 99.99% precision. The classification layer operates at an exceptional 96% accuracy, ensuring queries are understood with razor-sharp precision.

But we didn't stop there our RAGAS scores consistently range between 0.8 and 0.97, guaranteeing high-quality responses with deep contextual relevance. And our biggest triumph? An industry-leading average response latency of just 3.7 seconds, redefining real-time interaction and making AI feel faster, smarter, and more seamless than ever.

Shaping the Future: AI That Thinks, Learns, and Evolves

Schaeffler's AI-powered knowledge system has transformed efficiency, cutting troubleshooting time, reducing machine downtime, and ensuring seamless knowledge retention across its global workforce. Teams now collaborate effortlessly, accessing critical insights instantly, no matter where they are.

Looking ahead, we are building AI agents that go beyond answering questions they think, decide, and automate workflows for real-time, intelligent solutions. Powered by Azure OpenAI's reasoning models, these agents will make autonomous decisions, optimize tool usage, and benchmark results for continuous improvement.

Reinforcement learning will further refine responses by dynamically updating context based on user feedback. Long-term and short-term memory will ensure AI remembers past interactions, making conversations more fluid and intuitive. The future of AI at Schaeffler isn't just smart it's adaptive, automated, and always evolving.

Stay tuned for more AI success stories, solution deep dives, and expert insights.