



AI in the Factory and Beyond

7 High-Impact Use Cases in
Production and Diagnostics



Introduction

AI That Understands Your Data – and Your Business

Manufacturers and process industries generate **massive volumes of unstructured data** – from PDFs and manuals to logs, emails, and even images, video and audio recordings.

While GenAI holds promise, it struggles to reliably ground responses without the right context. That's where vector search comes in.

By combining GenAI with vector search and tools like MongoDB Atlas Search, organizations can build intelligent, context-aware systems that connect users to the information they need, fast.





“The challenge is no longer collecting data—it’s getting actionable insights from it. AI has the potential to transform manufacturing, but only if it’s grounded in the right context.”

– McKinsey AI in Manufacturing Report

1. Predictive Maintenance

Predictive Maintenance that Thinks Ahead

Unplanned downtime is one of the most expensive problems in manufacturing. While traditional sensor data can help spot anomalies, it rarely tells the full story. Valuable context – like maintenance records, technician notes, or operating manuals – often lives in unstructured formats that AI can't easily access.

By combining vector search with GenAI, organizations can connect structured sensor data with unstructured content. The result: a smarter system that doesn't just alert on anomalies – it understands them.

Imagine being able to query, "Has this issue happened before when the temperature spiked?" and instantly surface historical logs, maintenance tickets, and resolution steps – even if the terminology is different.

This is predictive maintenance that's not just reactive, but truly intelligent.



2. Knowledge Assistants for Engineers

Your Documentation, Searchable Like Google – But Smarter

Engineers and technicians often spend more time searching for information than solving problems. Critical knowledge is scattered across manuals, SOPs, SharePoint sites, and tribal notes – all written in different formats, by different people, over many years.

Traditional search tools rely on keywords, which often miss the mark. But with **vector search** and **GenAI**, you can give your team a smart assistant that actually understands what they're asking – and returns the most relevant answers from across your knowledge base.

Whether it's a question like *"How do I reset pressure valve 2A?"* or *"What's the torque spec for the actuator arm on line 3?"*, an AI-powered assistant can instantly surface the right paragraph, diagram, or SOP – even if the question doesn't use the exact wording.

It's like having your best technician available 24/7 – powered by your own data.



3. Root Cause Analysis (RCA)

Smarter and Faster Root Cause Analysis – Without the Guesswork

When production issues arise, finding the root cause is often more art than science. Teams sift through logs, operator notes, maintenance records, and sensor data – hoping to spot patterns based on experience or intuition.

But what if your system could do that for you?

With **vector search and GenAI**, manufacturers can automatically correlate new issues with historical incidents, even if they're described differently. AI can recognize similarities across unstructured data – log files, handwritten reports, emails – and surface relevant past cases, resolutions, or anomalies.

The result: faster RCA, fewer recurring problems, and less reliance on tribal knowledge.



4. Multimodal Diagnostics (Text + Audio)


Diagnosing Problems with More Than Just Data

Industrial equipment often *sounds* wrong before it *looks* wrong. But traditional diagnostic systems can't make sense of things like audio, operator notes, or irregular log entries.

By using **vector embeddings** and **multimodal AI**, manufacturers can build systems that understand patterns across different types of input – from machine noise to technician text notes. Just like in the demo where a system detected whether a yellow plastic card or a metal card touched a fan, these tools can now classify real-world sound anomalies and correlate them with known issues.

Even better, GenAI can explain why a warning was triggered, and what similar issues looked like in the past – across all your unstructured data.

This means earlier detection, smarter alerts, and fewer “I’ve never seen this before” moments on the floor.



5. Digital Twin Augmentation

Making Digital Twins Smarter with Real-World Context


Digital twins promise real-time visibility and simulation, but they often lack one critical ingredient: **context**. Most twins are built from structured data – sensors, CAD models, and control systems – but overlook the rich, messy knowledge stored in documents, logs, emails, and human input.

By integrating **vector search** with digital twin platforms, you can enrich your models with insights from manuals, maintenance reports, inspection logs, and more. This makes your digital twin not just a mirror of your machine – but a system that *understands* the machine.

Engineers can query the twin in natural language:

“What happened the last time this motor overheated for 3 hours?” – and get back relevant incident reports, past fixes, and recommended actions.

That’s not just simulation – that’s insight.



6. Quality Control + Defect Prevention


Finding Defects Before They Happen

Quality issues are often only discovered at the end of the production line – when it’s already too late. But in reality, the clues are usually there much earlier: buried in lab results, technician notes, test reports, or even photos and video.

With **AI-powered search and vector embeddings**, manufacturers can uncover subtle signals hidden in unstructured quality data. GenAI can help surface trends, correlate test outcomes with specific batches or settings, and even flag language in inspection notes that has historically led to failures.

It’s not just defect detection – it’s defect prevention.

The result: fewer surprises, lower scrap rates, and higher confidence in every unit that leaves your facility.




7. Supplier + Compliance Risk

Taming Supplier and Compliance Risk with AI

Supplier risk doesn't always show up in dashboards. It hides in contracts, certifications, audit reports, delivery logs, and email threads – the kinds of documents that traditional systems can't easily analyze.

With **vector search and GenAI**, companies can surface and track risks across all of this unstructured content. Want to know which suppliers have had recurring quality issues? Or who failed to meet compliance thresholds last year? Just ask – and the system can retrieve relevant documents, summarize findings, or alert you to patterns you might have missed.

You don't need a giant compliance team – just smarter tools that work across your data.





Actual Use Case


Business Challenge: Reducing Diagnostic Times in Automotive Aftersales

A Global Auto Manufacturer's Push to Modernize Diagnostics

Aftersales service is a critical revenue stream in the automotive industry. For one global automotive giant, diagnosing complex engine problems across a vast dealer network was slow, costly, and inconsistent. Traditional methods relied heavily on technician intuition and manual reviews of text notes and technical bulletins.

The company launched an ambitious project: to develop an **AI-powered diagnostic tool** that could analyze **engine sounds** and **text-based descriptions** to identify known issues and recommend solutions. The goal was to improve **repair speed**, **customer satisfaction**, and **revenue throughput**.

But after 9 months, the prototype was struggling to deliver accurate results – and the team needed a fresh approach.



AI for After Sales Automotive Diagnostics

- **Business need** : Identify similar incidents that already have a solution
- **Use Case** : Enable comparison between incidents and trace similar incidents that already have a solution.

Current State

Customer Complains



Cotech select a symptom


40 different symptoms



Cotech identify a solution among the possibilities



Cotech visualizes the solution Technical Bulletin (TB)



Future State

Customer Complains
Diagnostic comments

Search similar customer complaints or diagnostic comment symptom and directly push a solution

Cotech visualizes the Solution
Technical Bulletin (TB)


The Solution – A Phased, AI-Driven Approach with Pureinsights + MongoDB

From Setback to Success: Rethinking the Diagnostic Process with AI

Pureinsights joined the project to reshape the strategy. Instead of starting with complex sound analysis, the team began by mining **technician and customer text descriptions** using **vector search, semantic similarity, and Retrieval-Augmented Generation (RAG)** – all powered by MongoDB’s native document and vector stores.

- Once the text foundation was strong, the team advanced to **audio diagnostics**:
- Isolating engine sounds from recordings
- Turning sound into vectors
- Matching against a database of 200,000+ case files

The result is a scalable system that compares both **text and audio inputs** to known issues – dramatically reducing diagnostic time. The solution is already live in select dealerships, with broader rollout planned in phases throughout 2024.





“AI is no longer a futuristic concept – it’s a competitive edge. The companies that win are the ones that make innovation real, not just possible”

- Kamran Khan, CEO, Pureinsights

AI in the Factory and Beyond

Getting Started with Pureinsights

You don't need a massive AI transformation project to get started.

At Pureinsights, we help organizations in manufacturing and process industries take practical first steps with GenAI and vector search – using the data they already have.

Whether you're exploring predictive maintenance, AI search assistants, or multimodal diagnostics, we make it easy to move from idea to impact. Our prescription for success can include:

- ✓ Running a low-friction pilot on your own content and data
- 🔍 Ingesting and vectorizing manuals, logs, sound files, and more
- ⚙️ Integrating search tools like MongoDB Atlas Search, and AI platforms like OpenAI, Google Vertex AI, or Amazon Bedrock
- 🧠 Building intelligent copilots, assistants or agents tailored to your operations

Get started today: 📞 Book a [free consult](#) or [hands-on demo](#)





Additional Blogs, Demos and References

[AI-Driven Noise Analysis for Automotive Diagnostics | MongoDB](#)

[Automotive Diagnostics Using Atlas Vector Search | MongoDB](#)

[Demo: AI-Powered Search with Pureinsights Discovery and MongoDB Atlas](#)

[Demo: AI Powered Website Search with Pureinsights Discovery](#)

[Modernizing Search with Generative AI - Pureinsights](#)



About Pureinsights

Pureinsights transforms how organizations interact with information through intuitive, human-centered solutions. We deliver AI Search, Assistants, and Agents powered by Generative AI, Vector Search, NLP, and our Discovery Platform—helping businesses unlock insights and drive innovation across use cases from search to audio diagnostics.

Learn more at www.pureinsights.com or contact us at info@pureinsights.com



Consulting



Technology



SearchOps™

Copyright © 2025 Pureinsights. All rights reserved. Pureinsights™, its logo, Pureinsights Discovery™, and SearchOps™ are trademarks of Pureinsights Technology Corporation. Contents of this document may not be reproduced or used without express permission of Pureinsights, or without attribution to the source. This includes any derivative content submitted to generative AI for summaries or rewrites of the original content. This document is produced by consultants at Pureinsights as general guidance. It is not intended to provide specific advice on your circumstances. If you require advice or further details on any matters referred to, please contact Pureinsights.