PROGNOSTICS & HEALTH MANAGEMENT IN THE ERA OF AI

Introduction

What do airplanes, rockets, wind turbines, coffee machines, nuclear power plants, trains, hard disks, drones, bikes, bridges, ships — and basically any man-made system — have in common? They can, and eventually will, *fail*. If my coffee machine breaks, I shrug and make tea; but when I'm 10 km above the Atlantic, I'd much prefer the airplane to run smoothly. So, the real question is: ***How can we manage our systems effectively by anticipating failure without wasting valuable resources?***

This question has driven generations of engineers and scientists around the world to distil **Prognostics & Health Management (PHM):** a holistic paradigm that helps us monitor our systems' *heartbeat*, predict their performance, and plan courses of action — enabling the management of virtually *any* system.

Now turning our attention to the buzzword of the decade: **Artificial Intelligence (AI).** Far from being mere hype, AI is stepping in to close some of PHM's gaps. In this article, I examine two ideas: *Agents* which act autonomously; and *Large-Language Models (LLMs)* — capable of reading, reasoning and explaining — highlighting their potential to revolutionize PHM by giving our systems a *voice*, helping *us* to help *them*.

Why PHM Matters?

Maintenance engineers are familiar with reactive, scheduled, condition-based, preventive, and other maintenance policies — and attempt to choose the right one based on the problem at hand. Today, however, many industries are betting on two "hot" maintenance approaches: *predictive maintenance*, in which advanced models forecast likely scenarios so we can prepare in

advance; and *prescriptive maintenance*, which generates tailored action plans for precise, effective intervention.

Here, I use PHM as a unifying framework for delivering *health-based decision support* for maintenance, acting as a technical layer of asset management, embedding each system within its environmental, economic, and political context.

Thanks to decades of hardware advances — more powerful computers, expanded storage, and faster networks — we can move, process, and store massive volumes of data, unlocking unprecedented insights. In the past ten years alone, publicly available PHM datasets have more than doubled — a true game-changer that fuels the development of advanced predictive algorithms and flings the door wide open to AI.

The AI Revolution: The Potential of Agents & LLMs on Next-Generation PHM Systems

AI raises debates — philosophical, ethical, technical, and political — provoking constant controversy. I recall lying in bed, reading Max Tegmark's *Life 3.0* and picturing those dystopian and utopian scenarios, thinking, 'Will I ever live to see one of these?' I told myself, 'This won't happen anytime soon — maybe in a few decades.' How wrong I was!

With disruptive technologies like **GPT, DeepSeek, Gemini,** and other "cool kids on the block", we've transformed how we work — opening fresh opportunities in many fields, PHM included. I believe that one of PHM's greatest challenges is the lack of *cohesion mechanisms*: complex systems comprise numerous subsystems, operate under diverse conditions, and rely on expertise from multiple disciplines. Without cohesion, PHM setups can quickly fall apart — so how can AI help? *Let's dive into Agents and LLMs.*

Agent1: 'vibration spike on engine fan blade.' → Agent2: 'reserving hangar slot for blade inspection.'

Articles such as AI-2027 (https://ai-2027.com/) highlight the growing impact of software agents in modern society. These autonomous entities — products of **Agent-Based Modelling (ABM)** — can *perceive, act, and learn.* By embedding sensing, decision-making, and adaptation within each agent, ABM supports emergent, system-level behaviours that arise bottom-up from local interactions.

Conventional heuristics—valued for interpretability — struggle in dynamic, high-dimensional spaces and quickly become outdated. By "dropping" agents into carefully crafted digital worlds, we let them navigate vast mathematical landscapes and adapt policies through experience, refining strategies on the fly. ABM's modular design also supports *context-aware agents* that learn optimal strategies for different scenarios and *multi-agent systems* that collaborate toward shared goals.

Applications of ABM in PHM are booming. Agents can manage other models' hyperparameters (defined before model calibration) as a *control mechanism* — enabling diagnostic and prognostic algorithms to adapt continuously and maintain peak performance — or study component degradation profiles to *learn maintenance policies* that minimize costs while maximizing reliability. These adaptive approaches turn maintenance into a self-optimizing, feedback-driven process.

‘My diagnosis is that you’ve experienced a severe anxiety attack.’ — J.A.R.V.I.S

One of my favourite Marvel characters is J.A.R.V.I.S., Tony Stark’s AI — not just for its sharp humour, but because it embodies the PHM systems I envision. In this scene — System: *“Tony Stark”* → Task: *“Diagnostics”* → Result: *“Identified severe anxiety attack*”— the AI communicates a complex medical insight in plain language. That “voice” bridges system, hardware, software, and end users; you don’t need to be an engineer, scientist, or medic to understand it. The key challenge is that this next generation of PHM systems must mesh advanced, multidisciplinary techniques beneath a seamless *communication layer.*

Advanced LLMs already exist — GPT is one of the most palpable examples — and it’s only a matter of time before PHM systems adopt them. As specialized communication channels, LLMs can translate technical analysis into clear, actionable recommendations between maintenance engineers and technicians. When combined with agents, they enable *co-design mechanisms* — true human-agent collaboration — that power decision-support tools. These tools can propose maintenance strategies tailored to each situation, ensure compliance with regulations and resource constraints, and deliver guidance in language everyone understands, across any domain.

By harnessing agents’ feedback-driven learning loops, we can refine maintenance strategies while revolutionizing data collection. A major bottleneck in PHM is the scarcity of rich, real-world datasets, which hinders model training and validation. Integrating agents with LLMs could enable centralized platforms to gather metrics on maintenance effectiveness, component health indicators, and operational contexts — streamlining messy data processes, generating meaningful datasets, and fuelling continuous improvement and innovation of PHM algorithms.

It’s Not All Rosy: What Could Go Wrong?

Despite AI’s promise, we must balance rapid innovation with robust safety. In PHM, trustworthiness and explainability are crucial: if AI is to guide critical decisions, we need ways to verify and audit its reasoning, and to manage information overload as models evolve faster than can be monitored.

Agents bring their own hurdles: people distrust “black-box” actors, and tuning hyperparameters or ensuring ABM convergence in complex setups remains challenging. To deploy agents reliably in PHM, we need novel algorithms that handle these dynamics while staying transparent and predictable across varied operational conditions.

Traditional LLMs and agent architectures weren’t built for PHM, so we’ll have to adapt them — equipping them with domain knowledge, embedding physics-based constraints,

Despite AI's promise, we must balance rapid innovation with robust safety. In PHM, trustworthiness and explainability are crucial: if AI is to guide critical decisions, we need ways to verify and audit its reasoning...

and tailoring agents to sensor-driven feedback. This demands sustained, focused research into AI-driven PHM technologies.

We also face ethical, legal, and regulatory questions: who is responsible when an AI-driven maintenance decision fails? Addressing accountability and compliance frameworks will be vital as PHM systems gain autonomy. Data privacy and security are also important, as PHM platforms will collect sensitive operational data.

Despite these obstacles, I'm convinced the journey is worth it. Whether we ultimately succeed or hit roadblocks, exploring these technologies promises breakthroughs in system management and sustainability — and personally, I expect plenty of fun along the way!

Author:

dr. ir. Lisandro Jimenez
Data Scientist,
lisandro.jimenez@nlr.nl
Royal NLR - Netherlands Aerospace Centre