

AI in Predictive Maintenance Software: How It Works

Originally published at <https://www.iflexion.com/blog/predictive-maintenance-software>

Published: December 16, 2020 | By Martin Anderson



In both digital services and manufacturing, the modest profitability of the average delivery pipeline makes downtime expensive. In the energy industry, operation and maintenance costs for offshore wind turbines eat up 20-35% of all revenue for generated electricity¹, while in the oil and gas sector they represent somewhere between 15-70% of total production costs².

Therefore it's no surprise that these industries, as well as others, are keen to adopt more intelligent and informed maintenance approaches through predictive maintenance software and intelligent analytics with the help of [AI developers](#).

In this article, we'll take a look at some new developments related to this growing field, as well as some of the challenges in taking maintenance management out of the steam era and into the industry 4.0 age.

Approaches to Maintenance and Failure Handling

There are three essential approaches to industrial maintenance scheduling, the first two dating back some centuries:

Reactive Maintenance (RM)

Here, equipment is run to the point of failure and then fixed or replaced.

The advantage of RM is that the Remaining Useful Life (RUL) of all equipment is fully utilized, in theory. In fact, the disadvantages of RM make it a deprecated approach for most of the manufacturing infrastructure:

- RM requires the maintenance of expensive spare inventory, or else a reliable supplier (whose costs may rise over time, or who may discontinue essential components).
- Equipment could last longer with maintenance than through being exploited beyond repair.
- Costs can exceed those for preventative approaches.
- Breakdowns can cause unanticipated secondary damage through vibrations and other related conditions, increasing general costs.

Preventative Maintenance (PM)


Preventative maintenance seeks to anticipate wear and damage by scheduling overhauls based on time passed and informed by prior knowledge of tolerance limits (such as manufacturers' warranty recommendations for maximum continuous usage, load limits, and known records of fault incidence).

PM is perhaps best exemplified by motor vehicles' mileage-based service schedules and offers the greatest chance to minimize downtime. However, its conservative assumptions about component life can prove expensive:

- PM can lead to unnecessary repairs and shortened component life in some cases.
- It's often based on theoretical failure rates rather than known statistics.
- Budget constraints can extend maintenance intervals, leading to greater need for unanticipated and sometimes catastrophic replacement, with the accompanying downsides (see above) when equipment 'fails early'.

Predictive Maintenance (PdM, aka 'Conditioned-Based Maintenance')

PdM takes a wider view of the potential range of causative factors in industrial breakdowns, gathering contributing data with sensor types ranging from thermal imaging and ultrasonic through to vibration monitoring and [computer vision implementations](#).



Machine learning models can collate and interpret across these data streams to identify conditions where failure is more likely to occur, and to make more targeted predictions about when preventative measures should be taken. It's estimated that 50% of preventative maintenance tasks could be eliminated by moving to the PdM approach in industrial processes and machine/vehicle maintenance³.

However, it must also be considered that PdM not only has a higher cost of initial investment for the monitoring system and interpretive framework, but also requires high volumes of data. Additionally, changes in data type and the evolution of the system can require expensive machine learning consulting or in-house development — a scenario that's often only practicable for high-volume operators.

What's the difference between reactive, preventative and predictive maintenance?

Predictive Maintenance for Bearing Faults

On the factory floor, a bearing assembly typically comprises a ring that houses spherical metal balls designed to help machine components or speed them on their way to the next stage in manufacturing.

Since the general principles of bearing assemblies apply also to vehicles and other types of machine than factory installations, bearing analysis is a pivotal area of research in AI-based preventative maintenance.

Machine Learning Approaches to a Pre-Victorian Problem

Prior to recent advances in GPU-driven machine learning, a wide range of 'classical' machine learning solutions were trained on bearing assembly PdM over the years. These include artificial neural networks (ANNs)⁴, support vector machines (SVMs)⁵, principle component analysis (PCA)⁶ and k-NN classification⁷, among others.


These older methods depend on feature engineering⁸, where the shape of the equipment and rotation/movement rate is manually calculated against other user-specified input characteristics prior to machine analysis. This approach is quite brittle and often cannot be re-used or adapted to similar cases, making it labor-intensive and costly in itself.

Deep Learning for Bearing Fault Analysis


Newer deep learning techniques are able to observe a greater range of contributing factors that are hard to include in a feature engineering model. These include:

- The influence of heat⁹
- Variations in lubrication quality and flow¹⁰
- External vibrations, where multiple vibration sources can reach damaging frequencies¹¹ not anticipated in either individual source

Deep learning approaches incorporate automated feature engineering, provide transferable solutions, and, according to Google's most prominent contributing AI architect, Andrew Ng, offer an exponential improvement in failure prediction as the input data volume rises:



The first notable deep learning application for bearing PdM came from Ghent University in 2016¹² in the form of a convolutional neural network (CNN) with automated feature extraction and vibration analysis.



Neural network-based approaches to this challenge are also able to calculate 'imponderables' such as degradation of lubricating oil and other contributing factors that are difficult to model in traditional feature engineering.

As a further aid to research in this sub-field, dedicated data resources for bearing-based PdM have come up online in recent years, including the Case Western Reserve datasets¹³, the Mendeley vibration datasets¹⁴, and the bearing dataset in NASA's PCoE collection¹⁵.

PdM for the US Armed Forces

In a military context, equipment failure is catastrophic by default. Consequently, the armed forces in general have contributed greatly to research into preventative maintenance measures, and are committed to engagement with new AI-driven methodologies for ensuring the integrity of military vehicles.

The US Air Force announced in the summer of 2020 that it will greatly expand its current AI-based preventative maintenance operations¹⁶. According to the UAF's deputy chief of staff for logistics, engineering and force protection, this decision, estimated to cost \$3 billion, has been necessitated by the increase of state-based signals attacks around the world over the last decade — a factor set to make repairs and maintenance more difficult in future conflicts.


It's thought by some¹⁷ that the US Air Force has been inspired in this respect by the PdM innovations of Delta Airlines¹⁸ in the wake of a much-publicized technical outage in 2016¹⁹.

Across all sectors of the US military, issues around the 'right to repair'²⁰ cause additional logistical complexity. However, the problem extends well into the private sector, notably in agriculture²¹ and consumer vehicle maintenance²². Here PdM has much to offer, if it can overcome lobbyists' pressure²³ to maintain a more profitable 'reactive' status quo.


Predictive maintenance has much to offer in the US military if it can overcome lobbyists' pressure to maintain the 'reactive' status quo.

New Research for PdM across Industries


One 2020 paper²⁴ out of the Politecnico di Torino in Italy proposes a method of data mining to enable a regression-based PdM system that can predict vehicle faults across a range of possible data availability: a) vehicle-specific data, the ideal scenario; b) analogous data from similar types of vehicle; and c), a generalized vehicle model, which can be used when the vehicle is new, and no model-specific data yet exists.




Another Italian-led research initiative²⁵ offers a 'generalized' machine learning model for fault detection in centrifugal pumps for the oil and gas industries with approaches using support vector machines (SVMs) and multilayer perceptron (MLP).




The University of Bologna leads a recent paper²⁶ offering a fuzzy-rule-based approach to predictive maintenance for the Tier-1 data center supporting the Large Hadron Collider (LHC) at Geneva. The generalized methodology used by the study's Gaussian fuzzy classifier (eGFC) offers potential to transfer across to other similar data center and computing systems.



A Spanish-led paper²⁷ has proposed a real-time monitoring system for wind turbines wherein high volumes of historical data are fed into an always-on fault-tolerant network, with dashboard access.



Michigan State University recently published details of a methodology designed to develop 'universal diagnostics' across different types of vehicle by leveraging in-vehicle IoT sensors, as well as the media capabilities and GPU power of modern smartphones²⁸.



Class Imbalance in Training Predictive Maintenance Models

Unless they are carefully configured, machine learning systems tend to discount very rare events as irrelevant 'outlier' anomalies²⁹. Naturally, in the case of PdM, the identification and analysis of such rare (i.e. '*machine fault*') events is the entire objective!

This Class Imbalance can be addressed by oversampling the fault data³⁰, under-sampling the baseline data, or by down-weighting the classes used to compute the results³¹.

Synthetic sampling³² can be another useful approach, though its value can only be known for certain in practice rather than theory.


Security in Cloud-Based Predictive Maintenance Systems

In many cases, the costs of establishing a PdM framework can be lowered by using cloud-based services and open-source frameworks, typically applied in [enterprise-grade AI](#). The downside is that the attack surface for the system will be greatly increased — not a minor consideration when assaults on industrial frameworks such as Supervisory Control and Data Acquisition (SCADA) have grown in 2020³³.

Outright Denial of Service (DoS) attacks are damaging but usually lead to quick resolution and long-term countermeasures afterwards. A more effective type of incursion against PdM systems is a False Data Injection (FDI) attack³⁴.

In an FDI, the mechanisms of the targeted framework are used against themselves, either to cause unnecessary expense and downtime through premature maintenance requests or to permit equipment to fail beyond the threshold where the PdM system would normally have intervened. In critical cases, such an attack can threaten lives³⁵.

Typically, sensors are the weak link in the chain³⁶, since AI-based PdM systems are likely to use IoT-based information feeds from affordable or specialized sensors. The evolving state of long-range and near-signal network connectivity systems increases the opportunity to explore vulnerabilities in a growing number of network protocols:



Research out of the University of Missouri in 2019³⁷ tested the efficacy of three deep learning techniques to predict Remaining Useful Life (RUL) and to protect against FDI attacks: long short-term memory (LSTM); gated recurrent unit (GRU); and a convolutional neural network (CNN).

The project tested its system against a model of a turbofan engine in NASA C-MAPSS Aircraft Engine Simulator Data³⁸. GRU is a gating mechanism in a recurrent neural network that's capable of more refined training and higher specificity than an LSTM unit.

Other studies confirm^{39,40} that a GRU-based approach seems to be the way forward in defending industrial systems against FDI attacks, and generally requires less training time. However, GRU sacrifices a certain degree of accuracy to LSTM in favor of greater resilience against FDI incursions.

A GRU-based approach seems to be the way forward in defending industrial systems against FDI attacks.

The Future of AI-Based Preventative Maintenance

The use of machine learning in PdM systems is a developing field, currently lacking in national standards or even sector-based standardization, though there are initiatives and roadmaps that aim to at least set common terms and definitions⁴¹.

As expected, data gathering is a central issue among [machine learning challenges](#): NIST and car manufacturer Ford estimate⁴² that, unhelpfully, 98% of data gathered for PdM approaches are from healthy machine systems.

The Cost of PdM Implementation


There are logistical problems in developing 'test-to-destruction' datasets for rapidly evolving products: time-based endurance testing is costly and protracted, and in many product lifecycles the testing period would exceed the commercial availability of a component version. Further, manufacturers have little incentive to develop this information.

Since such metrics are more valuable for PdM systems than 'maximum load' ratings, it seems that the burden of time-based dataset development will fall on the end user. Passive monitoring systems based around computer vision, ambient vibration logging, and thermal imaging are more attractive here than the great expense involved in custom-fitting embedded sensors to equipment and components that have usually made no provision for this.

Generalizing the Principles of PdM

In terms of off-the-shelf products, predictive maintenance capabilities are found most often as a component of general asset management software⁴³. Though the specificity of manufacturing installations is an obstacle to a more generalized approach, new research promises the establishment of general principles in PdM, at least on a per-industry basis, particularly regarding PdM for semi-conductor RUL⁴⁴.

In any case, a clear rise in scientific papers centered on AI-based PdM over the last five years is a positive sign in this regard:



Afterword

In this sector, as elsewhere, sustainability is an intensely political topic^{45,46}. The growth of predictive maintenance is a potential threat to suppliers that have become financially dependent on after-market demand for parts and repairs⁴⁷ and that have already re-evaluated their core economic models, at some expense, to accommodate the 'disposable' culture that emerged in the late 20th century⁴⁸.

Though this incumbent burden is a challenge to the evolution of PdM, an ongoing shift in political and public thinking towards more ecological and economical manufacturing processes offers hope that the suppliers will increasingly adopt — and eventually champion — AI-based PdM.

In the meantime, the challenge remains to distil essential principles of industrial damage data into models that are transferable across systems and across a broader range of equipment, until the sector reaches a better level of commoditization, generalization and accessibility.

The shift to more ecological and economical manufacturing offers hope that the suppliers will increasingly adopt AI-based predictive maintenance.