Are Your Predictive Analytics Truly Predictive?

novity®

Novity

963 Industrial Road, Suite I
San Carlos, CA 94070

www.novity.us

Maintenance, repair and overhaul schedules can be optimized according to actual failure timelines.

Unplanned downtime is one of the most significant pain points for industrial manufacturers today, costing them an estimated $50 billion each year. The risk is even greater for process manufacturing, where a critical equipment failure could result in the loss of an entire batch, environmental hazards, or safety risks. The adoption of digital technologies, such as the industrial internet of things (IIoT), promises to mitigate these threats by forecasting equipment failures and catching faults before they lead to unscheduled shutdowns. However, in practice, several challenges arise when maintenance personnel and operations leaders work to implement an IIoT solution aimed at eliminating unplanned downtime.

As the various technical fields that support predictive maintenance (PdM) solutions have matured, the offerings and approaches available on the market have grown in scope and variety. Effectively sorting through these different solutions can become an effort in its own right, even before any implementation work has begun. Even for those early adopters who have been working on implementing IIoT solutions for years, there is often a disconnect between the expectation of what a solution will offer and the actual output of the product.

The narrative surrounding data analytics technologies, such as machine learning and artificial intelligence (ML/AI, here used interchangeably), is often the promise of a platform with predictive analytics that can predict when and how a piece of equipment is going to fail. In reality, the term "predictive" is misused for inherently nonpredictive technologies. Although nonpredictive technologies provide some value on their own, it should be clear what is truly predictive and what is not.

Unplanned downtime costs industrial manufacturers $50 billion annually.

Diagnostic Vs. Predictive Analytics

For an algorithm or software platform to be predictive, it should provide information on an event in advance of the actual occurrence of the event.

Currently, nearly all solutions advertised as predictive actually operate in a diagnostic fashion by providing explanatory insight into the current operation or condition of an asset or system. Diagnostic solutions take real-time sensor data and provide information on the current condition or performance of the monitored assets. Top-tier solutions can provide real-time notifications of minor problems that are known precursors to more significant problems, which provides value to the end user. However, this scenario is not predictive, as no information has been provided regarding the time or severity of any future events.

Online condition monitoring is a prevalent diagnostic technology that has seen recent advances and enhancements. Condition monitoring on its own only provides access to data. Online or continuous condition monitoring enhances this access by providing critical data in real-time over the complete duration of the asset's operating window. More sophisticated solutions may apply diagnostic analytics on top of condition monitoring to present the data in a form that is more easily interpreted by operators and maintenance managers.

With the advent of condition monitoring – particularly online condition monitoring – maintenance programs have shifted from schedule-based maintenance to condition-based maintenance (CBM). In schedule-based maintenance, maintenance is performed at regular intervals for a particular asset and its components. Typically, the schedule is derived from the original equipment manufacturer's recommendations. On the other hand, maintenance scheduling for a CBM program is derived from the assessed condition of the equipment. This allows maintenance managers more flexibility in when to schedule maintenance.

In contrast, predictive analytics paradigms, such as prognostics and predictive maintenance, go beyond the current state of an asset or piece of equipment and explicitly provide information on the time to failure. Predictive algorithms consider the current state of the machinery or process, as well as the loads and stressors, and predict the system's evolution. This provides additional insight into when and how the asset will fail. Similar methods can also be used to conduct "what-if" scenarios that predict the hypothetical outcomes of changes in the process, asset condition, or operation.

Are True Prognostics Possible?

Strictly speaking, a prognosis has three elements: time, location and severity of an event. However, this may seem difficult to accomplish in practice. Predicting when and how failures occur has long been considered a holy grail for industrial maintenance. As such, much research has been done on this subject. In recent years, this research has provided more accurate modeling, advancement of machine learning and data science and increases in computing power available to industrial operators. The application of such technologies to predicting machine failures has given rise to distinct areas of applied science, such as prognostics and health monitoring (PHM) and PdM. In fact, in the past five years or so, users have seen the commercialization of PHM and PdM research, with industrial end users beginning to realize value from this new technology.

By coupling accurate fault detection with prognosis, maintenance managers determine repair and maintenance time frames based on accurate estimations of the time to failure for an asset or component.

What makes prognosis possible is a confluence of several key technologies:

- New sensors specifically designed for use in IIoT settings that can operate in harsh or sensitive environments.
- Advances in machine learning and deep learning that make such approaches suitable for prognosis problems.
- Advances in physics of failure and simulation approaches that provide accurate predictive models of damage and failure progression.

When these technologies are combined with deep subject matter expertise, algorithms can be developed that take sensor data, track the progression of a specific asset's condition and provide a predictive model of when and how the asset will fail.

Deriving Value from Predictive Solutions

With diagnostic solutions, end users benefit from greater insight into the health of their assets and can realize value when faults are identified early in their progression. Such a value proposition is often referred to as actionable insight. It represents the new paradigm offered by typical implementations of IIoT, in contrast to the traditional practice of schedule-based maintenance. The success of diagnostic analytics also depends heavily on fault detection accuracy. With a high rate of false positives and false negatives, operators and managers have a difficult time judging which events should be acted on. Even when fault detection accuracy is high, maintenance managers cannot immediately address many early faults and still need to schedule service in advance. Without some sense of the time frame for a fault to progress to catastrophic failure of an asset, and without a reasonable estimate of the severity of such a failure, maintenance scheduling remains suboptimal, and unplanned downtime cannot be eliminated.

By coupling accurate fault detection with prognosis, maintenance managers determine repair and maintenance time frames based on accurate estimations of the time to failure for an asset or component. For instance, if a critical fault is detected, but failure is not expected for several months, more time can be allotted to prepare for repair and replacement activities so that disruption to production can be minimized or eliminated. Conversely, if a critical failure is predicted to occur in a matter of days, an immediate response can be prioritized with confidence.

The time horizon for prognosis is another aspect to consider. Emerging technologies are increasing the time before failure when a fault is detected by using metrics and features more sensitive to the mechanics of degradation and failure. With the expansion and development of the physics of failure and model-based prognostics techniques, algorithms can detect the earliest signals of incipient faults, often increasing fault detection lead times from days to weeks in advance of failure.

# Selecting a Predictive Analytics Platform

When selecting an IIoT solution, plant operators and maintenance managers should first define the expected outcomes of such an effort. If, after careful consideration, there are reasons to pursue an approach that provides failure predictions in addition to failure detection, then the following features for a predictive analytics platform should be considered when evaluating different solutions.

The most common approach to providing failure predictions is through a remaining useful life (RUL) estimation. Such an estimation will give an operator the amount of time before a failure is expected to occur on an asset based on the condition determined from sensor readings. Another parameter of interest is the uncertainty in the estimation, providing both a sense of estimation reliability and the timeframe within which failure is expected.

The prognosis algorithms should consider changes in asset operation or process conditions to produce the most robust and accurate RUL predictions. This can often be achieved using model-based or physics-based approaches. The RUL will then adapt to changes in process or operating point. Furthermore, by incorporating the physical processes of degradation and specific equipment characteristics, the features most sensitive to faults and degradation mechanisms can be selected, and accurate predictions can be made.

Methods also exist to provide information on the location and cause of failure. These diagnostic methods include fault detection, isolation and estimation. When these elements are combined, an operator or maintenance manager is provided with far more than an indication that something is wrong with an asset. Appropriate maintenance can be planned much further in advance and operations can be optimized to prevent unexpected failures completely.

With the expansion and development of the physics of failure and model-based prognostics techniques, algorithms can detect the earliest signals of incipient faults, often increasing fault detection lead times from days to weeks in advance of failure.

Case Study: Reciprocating Compressor Leakage Prediction

To illustrate the difference between diagnostics and prognostics, a reciprocating compressor operating on natural gas for a midstream oil and gas production company is examined. One of the most common failures on reciprocating compressors is the cylinder valves. These valves are actuated via pressure and experience high-impact stress continually during operation, making them susceptible to fatigue failures.

With a diagnostics approach, the most commonly implemented methods are threshold-based anomaly detection schemes. A relevant feature is tracked and its value is compared against some allowable threshold. The accuracy of such an approach relies on how well the threshold is chosen. Early detections can be achieved if the threshold is very close to the feature value during normal operating conditions. However, typically this also comes with additional false positives.

To reduce the false positives, the threshold may be adjusted to only trigger alerts when the deviations are large. This in turn results in alerts that are triggered much later when the fault is typically more severe.

A prognostics approach addresses these shortcomings by providing information on the trends in the tracked feature. Firstly, a feature must be chosen that is a strong indicator of the severity of the faults that the end user is interested in. With this condition met, trends in the feature tend to be noticeable long before thresholds are reached, including early-warning thresholds. Thus, a decrease in predicted RUL generally provides the earliest warning of an impending fault. Furthermore, since RUL estimations are computed from trends in the data, they are usually very robust to noise in the feature. As seen in image 1, the RUL has begun to decrease in response to the downward trend in the data, long before the damage level has deviated significantly from the healthy condition.


IMAGE 1: The end of life of a reciprocating compressor valve is predicted over a month in advance of the actual fault.

Often, machinery faults fail gradually at first, then very rapidly near the very end of life. Thus, once the severity of a fault has reached a level where alerts can be confidently triggered, it's typically very close to the component's end of life. As illustrated in image 2, warning thresholds for this fault triggered alerts repeatedly once the fault severity reached the warning threshold condition (red horizontal line). The earliest alerts may have been construed as false positives, since the damage is still not very significant. However, as seen in image 3, the later warning alerts occurred only a few days before the failure, while the prognoser accurately predicted that the end of life was only a few days away. As time progresses and the prognostics algorithm consumes more history of the growing fault, the estimated failure date becomes more stable and the prediction more accurate.


IMAGE 2: Later progression of a reciprocating compressor valve fault. Times when warning alerts are triggered are depicted with yellow dots.


IMAGE 3: Later progression of a reciprocating compressor valve fault. Times when warning alerts are triggered are depicted with yellow dots. Critical alerts are depicted with red dots.

Case Study: Prognosis of a Heat Exchanger

The value of prognosis can be further illustrated with the gradual degradation of a heat exchanger through fouling. Fouling is virtually every heat exchanger's most common failure mode. However, in many applications, the physical processes that drive fouling are only sometimes linear or easy to predict. For instance, some heat exchanger applications in the chemical processing industry exhibit self-cleaning behavior from time to time. Thus, the level of fouling can repeatedly fluctuate, resulting in false positives.

By including models of the evolution of the physical processes involved in fouling and the performance characteristics of the heat exchanger, changes in the operating conditions can be accounted for, and the non-linear behavior of the asset over time can be more accurately predicted. In Image 3, the end of life of a spiral heat exchanger is predicted well before substantial degradation has occurred. Examining the health history in the first figure only, a maintenance manager may conclude that cleaning the heat exchanger will not be necessary for many years. However, with the RUL prediction incorporating model-based prognostics, a more accurate estimation of the maintenance needs shows that cleaning should occur much sooner than would be considered with condition monitoring alone.

Although the prediction in the first figure has some amount of uncertainty and error associated with it, a robust prognostics algorithm decreases in both error and uncertainty over time as more history is available to the algorithm. Such a reduction in error and uncertainty can be seen in Image 4 as the health has reduced further. The signals characteristic of the failure grow as the degradation progresses and the prognosis algorithm becomes more accurate.
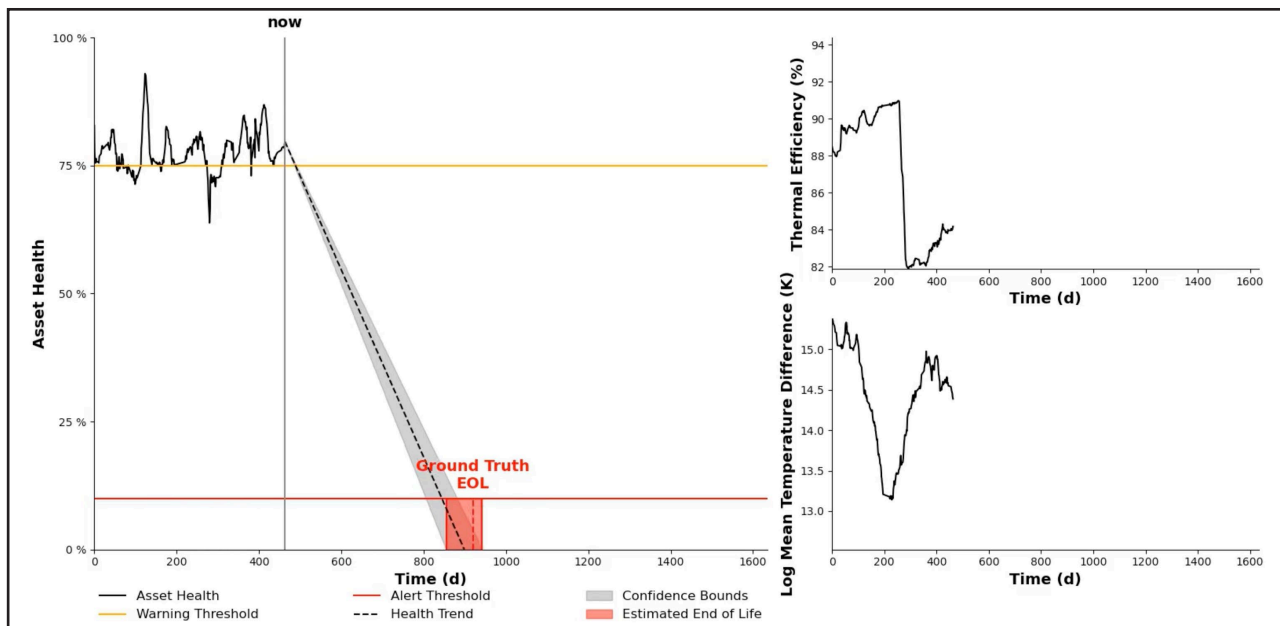

IMAGE 4: The end of life of a spiral heat exchanger is predicted well before substantial degradation has occurred.
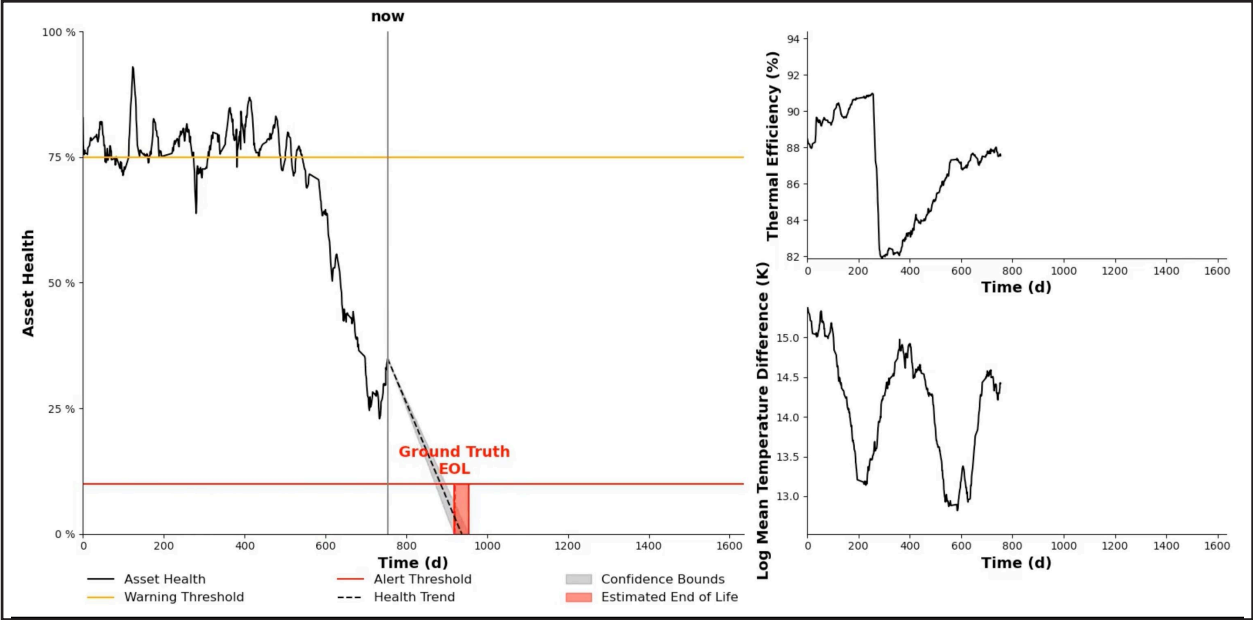
IMAGE 5: Reduction in error and uncertainty.

Process industry equipment users can drive significant improvements to their bottom line by informed application of predictive and diagnostic technologies. With new sources of data coming from sensors, and effective analytical tools to process these data streams, accurate and timely notifications of equipment problems can be provided, and appropriate action can be taken before any additional damage to mission-critical hardware occurs. Furthermore, by leveraging emerging prognostics technologies, maintenance personnel no longer have to guess when equipment is going to fail. Maintenance, repair and overhaul schedules can be optimized according to actual failure timelines, and the possibility of eliminating unplanned downtime can be realized.

novity®

Contact Us

Address:
963 Industrial Road, Suite I
San Carlos, CA 94070

Telephone: 1.833.375.4946

Website: novity.us