
ARUNDO

Predictive Equipment Maintenance

Anomaly Detection
for Industrial
Operations



Houston, USA

757 N Eldridge Pkwy Suite 800H
Houston, TX 77079

Oslo, Norway

Wergelandsveien 7
0167 Oslo

Stockholm, Sweden

Tomtebgatan 20
11338 Stockholm

Heidelberg, Germany

Mittlerer Rainweg 63
69118 Heidelberg

CONTENTS

Summary	03
The Changing Data Landscape Is Transforming Industrial Equipment	05
Predictive Equipment Maintenance In Theory And Practice	07
From Single-sensor Threshold Alerting To Anomaly Detection	09
Anomaly Detection Drives Strong Results Even With Limited Historical Event Data	12
Putting Theory Into Practice	14

AUTHORS

Lars Bjålie, PhD, Senior Data Scientist at Arundo Analytics. At Arundo, he has worked primarily on process optimization in the chemical industry. He previously held a position as a consultant at BearingPoint, where he worked as the lead data scientist within predictive maintenance.

Emre Kucuk, PhD, Data Scientist at Arundo Analytics. He works on developing and implementing industrial data science applications, which encompass anomaly detection for equipment monitoring. He previously worked on predictive maintenance projects at AVL and Fujikura.

Ellie Dobson, PhD, former VP Data Science at Arundo Analytics. She worked extensively on research and development of industrial data science applications, including anomaly detection for equipment monitoring. Ellie previously worked at Pivotal Software, The Mathworks, and CERN.

Amitav Misra former VP Marketing at Arundo Analytics. He worked closely with industrial equipment manufacturers and major asset owners seeking equipment monitoring solutions. He previously worked at EnerNOC Inc. and McKinsey & Co.

SUMMARY

Predictive equipment maintenance is rarely, if ever, an “out of the box” solution, but rather an iterative development process.

Achieving this goal requires a deeper understanding of the underlying data science and technology journey related to equipment monitoring and predictive analytics.

Heavy industrial companies are increasingly taking advantage of the Industrial Internet of Things and deploying machine learning applications in order to improve revenue, reduce costs, and improve safety across their operations. A common focus for these companies is to integrate predictive analytics - automated actions for certain highly likely outcomes - into their business processes. In particular, predictive maintenance for critical equipment is a key goal.

True predictive equipment maintenance involves automated notification of the specific mode of pending equipment failure: a message or alert to a human operator or control process that enables scheduled maintenance of the equipment at the lowest possible cost in terms of materials, labor hours, and equipment downtime.

In order for a machine learning model to produce accurate predictions to support this process, the model must learn

- (a) the universe of potential failure modes, and
- (b) how patterns of sensor signals indicate specific failure modes.

Despite the best efforts of data scientists and technologists, it is almost impossible to implement true predictive maintenance models without significant historical operating data related to both sensor data and a large number of repeated, specific failures, which is seldom available at the start of a company's predictive maintenance journey.

Therefore, we recommend a "roadmap" approach, starting with streaming data capture and threshold-based alerts, and moving quickly into advanced anomaly detection as groundwork for the ultimate goal of a true predictive maintenance system.



THE CHANGING DATA LANDSCAPE IS TRANSFORMING INDUSTRIAL EQUIPMENT

Leading companies in heavy industries – operators and suppliers in energy, maritime, utilities, chemicals, and other capital-intensive operations – are reshaping their approach to operating performance in response to the convergence of several long-term technology trends:

- Sensors continue to decline in cost and physical footprint.
- Sensor, device, and asset-level connectivity continue to improve in quality and cost.
- Data platform technology in the cloud enables rapid start-up and scale-up, and access to new and continuously developed tools.
- Machine learning tools and techniques are increasingly accessible.

This intersection of device connectivity and compute capability is often referred to as the Industrial Internet of Things (IIoT). For industrial companies, the ability to interact with physical equipment has never been greater. Companies now have the ability to access and analyze a previously unimaginable amount of data, arriving almost constantly from a variety of sources, and to make meaningful business decisions from this data continuously.

Industry leaders are taking advantage of this ability to increase revenue, decrease costs, and to create new business models.

Already, many asset owners and operators are “sensing up” their physical operations – even before finalizing the new business strategies, operating processes, and software tools required to realize value from new digital assets and data streams.

However, most companies are just beginning their digital transformation – they typically access, analyze, and make decisions based on just a tiny fraction of the potential data generated from their assets and equipment.

There are many ways to derive value from the Industrial Internet of Things, such as production process optimization, enhanced quality control, supply chain optimization, energy management and predictive maintenance. In this whitepaper, we drill down into the predictive maintenance theme, which has been shown to reduce equipment downtime and maintenance costs typically in the range of 15-20%^{1,2} by monitoring the condition of machinery and equipment in real-time, anticipating equipment failures so that maintenance can be scheduled proactively.

Sources:

1 <https://www.bearingpoint.com/en-no/insights-events/insights/predictive-maintenance-study-2021/>

2 <https://www.mckinsey.com/capabilities/operations/our-insights/the-future-of-maintenance-for-distributed-fixed-assets>

PREDICTIVE EQUIPMENT MAINTENANCE IN THEORY AND PRACTICE

In true predictive maintenance applications, a piece of equipment, such as a pump, compressor, or heat exchanger, alerts human operators or other systems about a specific impending failure mode in time for an intervention. In order to achieve this goal, “supervised” machine learning models (algorithms that learn to recognize labeled data patterns) can be trained on representative historical data containing several (correctly) labeled examples of all different failure event types.

However, such datasets are rarely available in actual industrial operations. For most field-installed equipment, data is commonly available in one of the following scenarios:

1. There is no historical sensor data, since we are in the early stages of setting up infrastructure to collect and store sensor data.
2. There is historical sensor data, but it is not labeled with historical failure events.
3. There is historical sensor data, and there is historical failure data, and with manual or automated pre-processing, the datasets can be joined to get labeled sensor data. However the failure events are relatively rare, and don't fully represent the entire universe of potential failures in sufficient numbers to train models for automated predictions.

These common data situations make the immediate application of true predictive maintenance systems quite rare for most heavy industrial equipment, since the quality and accuracy of machine learning model outputs are largely driven by the availability of historical failure data. Even though the absolute number of failure events needed to train a highly accurate model can be quite low (typically in the hundreds) if the failure differs enough from normal operations, without a sufficient number of historical failures, applying even the most sophisticated machine learning techniques is a futile effort.

To get around this data hurdle and get started with predictive maintenance, a common initial approach is to create threshold-based alerts for individual sensors (for instance, if temperature or vibration fall above or below specified levels, an automated notification is sent to certain users).

Unfortunately, industrial equipment typically exhibits a range of complex behavior, which makes it challenging to make accurate threshold-based alerts. They may fail to alert of impending failures, and also raise unnecessary alarms when the equipment is simply in a corner condition of normal operations, or perhaps in a different operational mode such as a ramp-up phase. False alarms are a significant challenge for equipment monitoring systems in heavy industry, as they often result in operators losing trust in the systems and ignoring alarms altogether, and failing to act in advance of major equipment failure

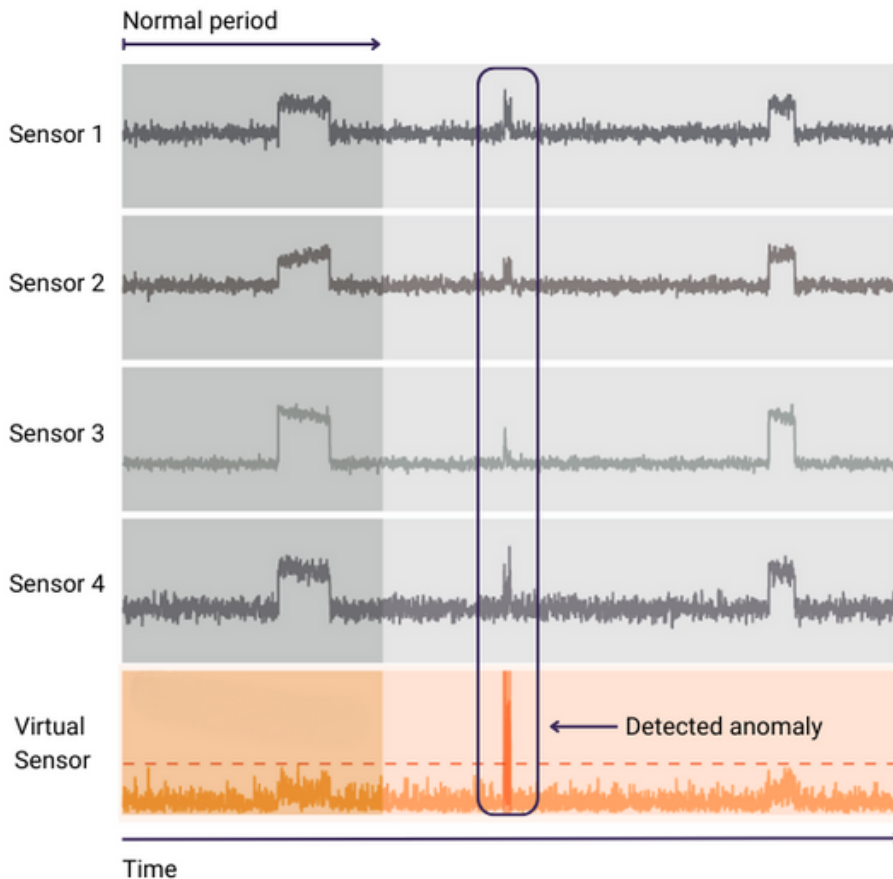
FROM SINGLE-SENSOR THRESHOLD ALERTING TO ANOMALY DETECTION

Moving beyond threshold-based models, anomaly detection, an “unsupervised” machine learning approach (algorithms that detect patterns in unlabeled data), can be used without the same stringent requirements on labeled failure modes as required by supervised models. These types of algorithms can often find more failures and produce fewer false alarms than single-sensor threshold alerts.

Arundo has several years of experience developing anomaly detection models on numerous pieces of industrial equipment. One of our standard approaches is to use time periods where we know the equipment operated normally (e.g. right after maintenance, or when it was new) to train an algorithm to learn what constitutes normal behavior across all sensors - in contrast to failure labels, this “normal period” data is very often available.


We can visualize the health of the equipment via a “virtual sensor” outputted by the algorithm, and choose a threshold for flagging anomalies based on the maximum virtual sensor value in the normal period. The trained model can then raise alarms when both previously unseen or known faults occur, producing more relevant findings than threshold-based alerts: It will not flag normal operating modes as anomalous, provided that they are represented in the time periods of normal operation the model was trained on.

This is illustrated in the figure below, where the equipment has two different operating modes, as seen in the normal period used to train the algorithm. For the remaining time, the algorithm (an autoencoder neural network built from a few fully connected layers) does not flag the same operating mode as an anomaly, but clearly identifies an anomalous pattern which can resemble the other operating mode to the naked eye.




One drawback of anomaly detection models is that they are not able to specify exactly which failure mode is going to occur - domain expertise is still needed to diagnose the exact failure mode (e.g. inspect the sensor data, the equipment itself, and the maintenance logs). However, such expert-supported diagnosis can in turn reveal insights that can be codified into threshold-based alerts that account for various operating modes, either by themselves or as a post-processing step from the outputs of the anomaly detector. This forms a two-step approach that can be repeated and deliver a handful of prescriptive alerts that are intuitive to the asset operators. Arundo has several years of experience applying such an approach, so for many processes and equipment we have existing prescriptive solutions which can be deployed from day 1.

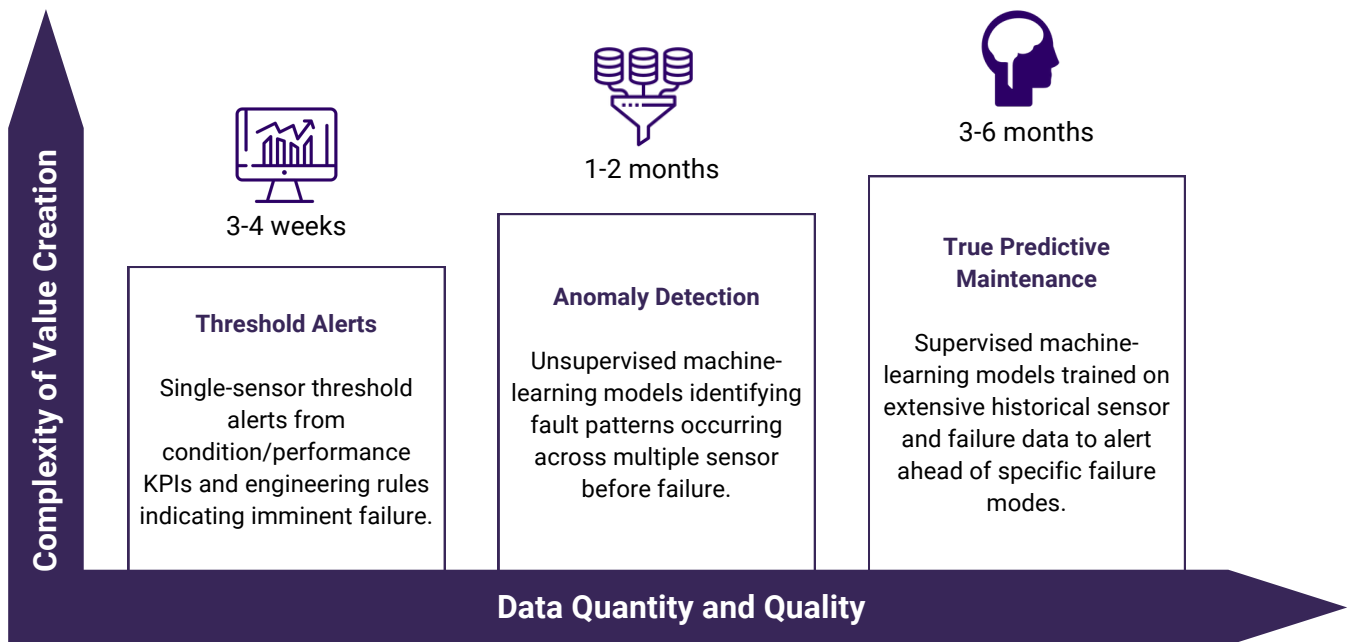
1. Train a general anomaly detection model with unsupervised learning




2. Build multiple specific models with actionable outputs from the anomalies




All in all, we recommend a “roadmap” approach, starting with streaming data capture and threshold-based alerts, and moving quickly into anomaly detection as a groundwork both for improved threshold-based alerts and the ultimate goal of a true predictive maintenance system. With such models running in production, anomaly types indicating different future failure modes can be monitored to build up a data set with labels, allowing the maturation into “true” predictive maintenance.



ANOMALY DETECTION DRIVES STRONG RESULTS EVEN WITH LIMITED HISTORICAL EVENT DATA



At a major national oil company, we used anomaly detection to identify potential compressor failures with adequate operational lead time for inspection, maintenance, or repair. The anomaly detection approach is now integrated into the company's equipment monitoring system.



A hydroelectric power plant exhibited complex operating behaviors. These included varying operating conditions in ramp-up, normal operation, and ramp-down modes. We deployed a density-based clustering model for anomaly detection that provides higher confidence estimates of potential failure than any physics-based model, engineering simulation, or equipment manufacturer estimates.



At a major offshore oil & gas producer, we provided anomaly detection as an integrated equipment condition monitoring application for heat exchangers and compressors. This is delivered through an application with alert functionality and ability for equipment specialists to access all information related to failures (sensor data, technical drawings, etc.) in one integrated view for rapid diagnosis to take preventive actions.

In aggregate, these applications of anomaly detection have the potential to save tens of millions of dollars across the three companies.



PUTTING THEORY INTO PRACTICE: MARATHON

From our experience working with mature industry partners, Arundo has developed Marathon, a comprehensive Asset Reliability Management software suite to put data-science models such as the anomaly detectors discussed here into production and drive maintenance decisions.

You can read more about Marathon here:

<https://www.arundo.com/marathon>



ABOUT ARUNDO

Arundo Analytics helps equipment operators to make better decisions and improve efficiency, reliability, and sustainability.

Our products link data science, cloud computing, and AI with industrial IoT, providing prescriptive actions beyond simple alerts. Broad industry experience makes us a powerful partner for digitalization and process optimization.

Maximize your business' capabilities with Arundo products.

Benefit from more than 250.000 development hours for:

- increased production levels,
- minimized downtime,
- enhanced reliability, and
- significant cost savings for your operations.

**Get the answers you need:
contact@arundo.com**

