



Xen.AI PdM (Predictive Maintenance) is an Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning based solution for predictive maintenance.

Overview

Predictive maintenance (PdM) is maintenance that monitors performance and condition of equipment during normal operation to reduce a likelihood of failures. Also known as condition-based maintenance, predictive maintenance has been utilized in the industrial world since the 1990s.

Predictive maintenance evaluates the condition of equipment by performing periodic (offline) or continuous (online) equipment condition monitoring. The ultimate goal of the approach is an ability to predict when equipment failure could occur (based on certain factors), and to perform maintenance at a scheduled point in time prior to this, i.e. when the maintenance activity is most cost-effective and before the equipment loses performance (within a threshold).

This results in a reduction in unplanned downtime costs because of a failure where, for instance, costs can be up to hundreds of thousands per day depending on industry. This is in contrast to time- and/or operation count-based maintenance, where a piece of equipment gets maintained whether it needs it or not. Time-based maintenance is labor intensive, ineffective in identifying problems that develop between scheduled inspections, and so is not cost-effective. So, the fundamental idea is to transform the traditional “fail and fix” maintenance practice to a “predict and prevent” approach. Predictive maintenance differs from preventive maintenance because it relies on the actual condition of equipment, rather than average or expected life statistics, to predict when maintenance will be required.

Some of the main components that are necessary for implementing predictive maintenance are data collection and preprocessing, early fault detection, time to failure prediction, maintenance scheduling and resource optimization through "just-in-time" in manufacturing.

Modern technologies based on integration of Artificial Intelligence (AI) and Machine Learning (ML), allow us to solve most of such problems in an effective way. They include learning from time series of historical data with known healthy periods of operation and also time points when certain failures occurred. Using this data one can train models that would be able to detect anomalies in data and predict probability of a failure.



Examples

Some typical examples of the problems include (a) monitoring of moving truck with sensors installed on engine, transmission, brakes, tires, and body parts, where we would measure such quantities as temperature, pressure, speed, torque, acceleration, GPS location and predict when some parts may start failing; (b) detecting false clicks or orders on an E-commerce website that would be characterized e.g. by a high activity from particular geographical locations, devices, and days of week, or multiple orders of same item; (c) detecting fraudulent credit card activities, (d) prediction of patient state in a hospital/ICU, (e) churn prediction (survival analysis).

Approaches

In spite of seemingly different problems across different industries, their solutions are based on quite similar approaches from the point of view of ML&AI.

1. Online condition monitoring and anomaly detection in time series data.

In this case, a typical problem setting includes a number of numeric and categorical variables changing over a time. One needs to detect a deviation from a normal behavior. It is important to understand here that most of the variables may correlate to each other, and thus we have to look for anomalies at a whole set of variables simultaneously rather than at individual ones.

We build a model using data for a “healthy” periods of activity that “knows” how healthy system should look like. Then based on a given vector of variables, $[v1, v2, v3, \dots, vN]$, for a given timestamp (or aggregated time unit, e.g. an hour), one can detect how big is deviation from a normal behavior, i.e. calculate an anomaly score. We relate this score with the input vector, and report a problem (see Fig. 1).




Fig. 1: Detecting anomaly in a time series data. Anomaly threshold is shown by the red line.

2. Failure prediction for a time in future.

Sometimes we would like to know when failures that happened in the past may happen again in the future. For this purpose, we train a model using a set of variables that change over time (time series data), and predict their behavior for a future.

If actual observed data deviate from the predicted trend by a significant amount (i.e. if it goes beyond estimated confidence interval), we report anomaly. Please see a typical plot below at Fig. 2 that describes this.




Fig. 2: Predicted and actual data versus time. Shadow area shows a confidence intervals on predicted variations. Red points in data indicate detected anomalies.

3. Time to failure

Provided that we have data on the operation lifetime of a certain item (e.g. engine), one can predict remaining working cycles (days) to a next failure. It is also known in the field as prediction of Remaining Useful Time (RUL). It can be done by constant monitoring of behavior of different sensors. At certain point of time, their simultaneous trends (e.g. measures in two sensors below at Fig. 3 started growing) may indicate an approaching failure. One can predict a confidence interval for such a failure time.




Fig. 3: Sensor measurement versus time.




Fig. 4: Training predictive models that can estimate remaining useful life and provide confidence intervals associated with the prediction.

4. Other applications

Predictive repair has a wide range of applications. Among them are finding patterns and classification in telematics and IoT data, E-commerce, finance and healthcare. Below we show application to recognition of heart diseases using ECG data, please see Fig.5.



Goals of Xen.AI PdM Solution

1. Together with a client, discuss available data, operational regimes and types of failures that should be addressed.
 2. Determine most influential variables related to each type of failure.
 3. Develop Artificial Intelligence & Machine Learning based integrated solutions to detect anomalies in data streams. Optimize anomaly thresholds to better balance between false positive and false negative rates.
 4. Predict future failures.
 5. Run models using online data streams.
 6. Develop a dashboard that would help track sensor behavior and failure status.
-

Xen.AI PdM Solution Overview

1. Feature analysis.

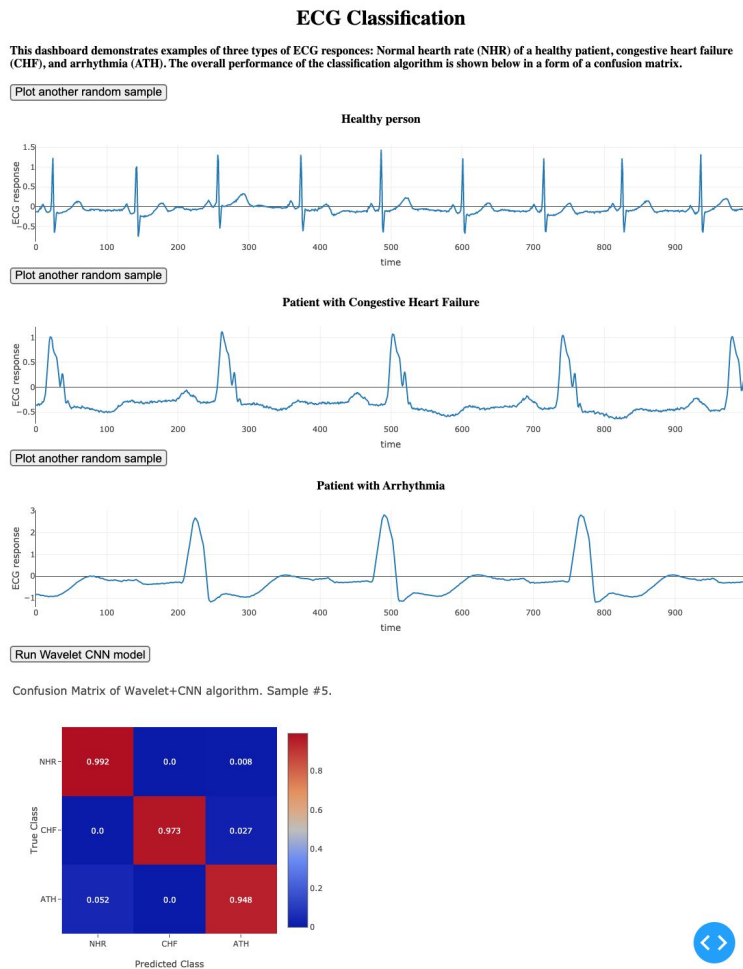


Fig.5: Classification heart problems using ECG data.

When considering data for a predictive maintenance system, we have to analyze data early to understand which features are important and which may be redundant. It would allow us to see which variables actually affect failures. Also, depending on where data is stored, it can be expensive to keep an excessive amount of data that is not going to be used. It also reduces time for model training.

2. *Understand the available data.*

Failure data might not be present, but operations data might show trends about how a machine degrades over time. Looking at the raw sensor data from a component, system, or machine with dozens or hundreds of sensors can be intimidating. Statistical techniques such as principal component analysis (PCA) can help reduce the dimensionality of such datasets and provide valuable insight into how equipment operates over time. Depending on what sensors are available, certain types of failures may require looking at several sensors simultaneously to identify undesirable behavior. Unsupervised learning techniques transform raw sensor data into a lower-dimensional representation, which can be visualized and analyzed much more easily than the high-dimensional raw data (see Fig. 5). Keeping the number of variables down to the minimum needed is also important for an accurate and more transparent model.




Fig. 6: Using principal component analysis to visualize how equipment trends prior to failure.

3. *Generate more failure data.*

Failure data is a crucial part of teaching algorithms to recognize the warning signs to trigger just-in-time maintenance. However, failure data may not exist if maintenance is performed so often that no or just few failures have occurred. To stop this from becoming a fatal deficiency, one needs to simulate failure data from a physical model of the

machine to supplement normal usage, varying parameter values, different system dynamics, or signal faults.

4. *Start small and gain confidence.*

Our philosophy: rather than trying to cover all possible failures, we choose a project using a deeply understood system (maybe with a high business impact) to start from. Make sure we understand the features and factors that affect the performance of the system, and build a predictive maintenance algorithm using simple models first (even using just linear and logistic regressions). Once it works and all trends and dependencies are clear, we apply that knowledge to more complex systems.

Second, when predictive maintenance algorithms begin to show promising results, we use current and historical data to test and validate models before moving to production. We use the domain knowledge within a client's team to tune models to predict different outcomes based on the cost/severity of those outcomes. To further validate models, add generated failure data similar to known historical conditions and test the system. This validation step will build confidence that the process is properly working.

5. *Usage of Open Source tools.*

To minimize dependence on vendors, we prefer to use open source tools like Python, Tensorflow, Keras, scikit-learn and other machine learning libraries.

6. *Parallel processing to speed up.*

Depending on frequency of model training and data volumes, we can utilize cloud computing cluster resources following the pay-as-go concept by paying for just execution time.

Building workflow for real applications

The predictive maintenance workflow typically includes following main 5 elements (please see Fig. 7 below):

1) Access sensor data.

Make sure the data is in the right format. Large data sets may not fit into computer memory, and will require out-of-memory processing techniques or a cluster.

2) Preprocess data.

Data in the real world is rarely perfect; it may have obvious outliers and noise that need to be removed to get a realistic picture of normal behavior. If the data has come from different sources, it will also need to be combined.

3) Extract features.

Instead of feeding sensor data directly into machine learning models, it is common to extract features from the sensor data. These features capture higher-level information in the sensor data. An iterative approach—in which features are added, new models are trained, and their performance is compared can work well here to determine the effectiveness of different features on the results.

4) Train the model.

In this step, we classify data as healthy/faulty, set thresholds for healthy/warning/failure states, and estimate remaining useful life for components. We'll need to create a comprehensive list of failure scenarios to predict, choose classification methods, and simulate models.

5) Deploy the model.

Generate code and deploy models as an application integrated with other applications in an IT environment. (Models may be deployed to embedded devices by converting them to a low-level language such as C.)




Fig. 7: Basic workflow for predictive maintenance.

Key Benefits of Xen.AI PdM Solution

The suggested solution can be helpful in a wide variety of industry fields, for example, Automotive, Railway, Manufacturing, Oil and Gas, Finance, E-commerce, Advertising. Possible positive outcomes include:

- Ability to monitor the failure prediction online and plan all related actions in advance.



Innovative Artificial Intelligence Solutions

Web: www.xen.ai Email: support@xen.ai

- Make very specific part replacement rather than a costly bulk screening of all potential causes.
 - Timely detection of failures can help reduce amount of unplanned downtime repairs and relevant costs.
 - Increase customer retention.
 - Increase company revenue.
 - Increase the company's brand value.
-

Contact Us

Web: www.Xen.ai

Email: support@xen.ai

USA:

Param Namboodiri

501, Gibson Dr, #2624

Roseville, California - 95678, USA

Phone: +1 408 221 6976

INDIA:

Shanawaz Hakeem

ES 11, Heavenly Plaza, Kakkanad

Kochi – 682021, Kerala, India

Phone: +91 907 488 7447
