# Understanding Wind Turbine Condition Monitoring Systems

More and more, the wind industry is recognizing the value of condition monitoring. This shift is driven by two major factors: many wind turbine warranties are expiring—exposing owners to the true operations and maintenance costs of wind farms, and component failures are driving excessive maintenance costs. To combat this, owners and operators are deploying condition monitoring systems (CMS) to detect faults before they cause secondary damage. Through this early detection, repair costs can be reduced, representing significant savings.

As interest has grown, so has the number of vendors offering systems. Almost every turbine OEM offers a CMS, and gearbox and bearings suppliers are also joining this crowded space. With several technologies available (e.g. vibration, oil debris, SCADA), determining which one provides the highest value can feel like comparing apples to oranges.

Unfortunately, there is no performance standard or benchmark for comparison. A prospective buyer is left with the difficult task of deciding what CMS will provide the right performance at the right price. To succeed in this assessment, it is important to focus on the overarching goal for the system: CMS are designed to provide users with recommendations that enable them to make optimal operations and maintenance decisions.

It is also crucial to understand the *process* through which a CMS converts a physical measurement (e.g. vibration, oil debris, temperature, pressure) into a recommendation for action. Though methods differ, there is a generic six-step process that all CMS follow:

1. **Data acquisition:** translation of the physical phenomenon into an analog measurement, which is then converted into digital format;
2. **Data processing:** conversion of the digitized measurements into meaningful indications of component health;
3. **Detection:** classification of the condition indicators as "normal" or "abnormal";
4. **Diagnosis:** validation of the fault and determination of its location and severity;
5. **Prognosis:** estimation of how much longer the faulted component will last before it needs to be replaced; and
6. **Recommendation:** determination of what maintenance action is necessary and when it should be performed.

Understanding how a particular CMS performs each of these steps gives a prospective buyer a much clearer understanding of the system's capability and the cost of each step in the process.

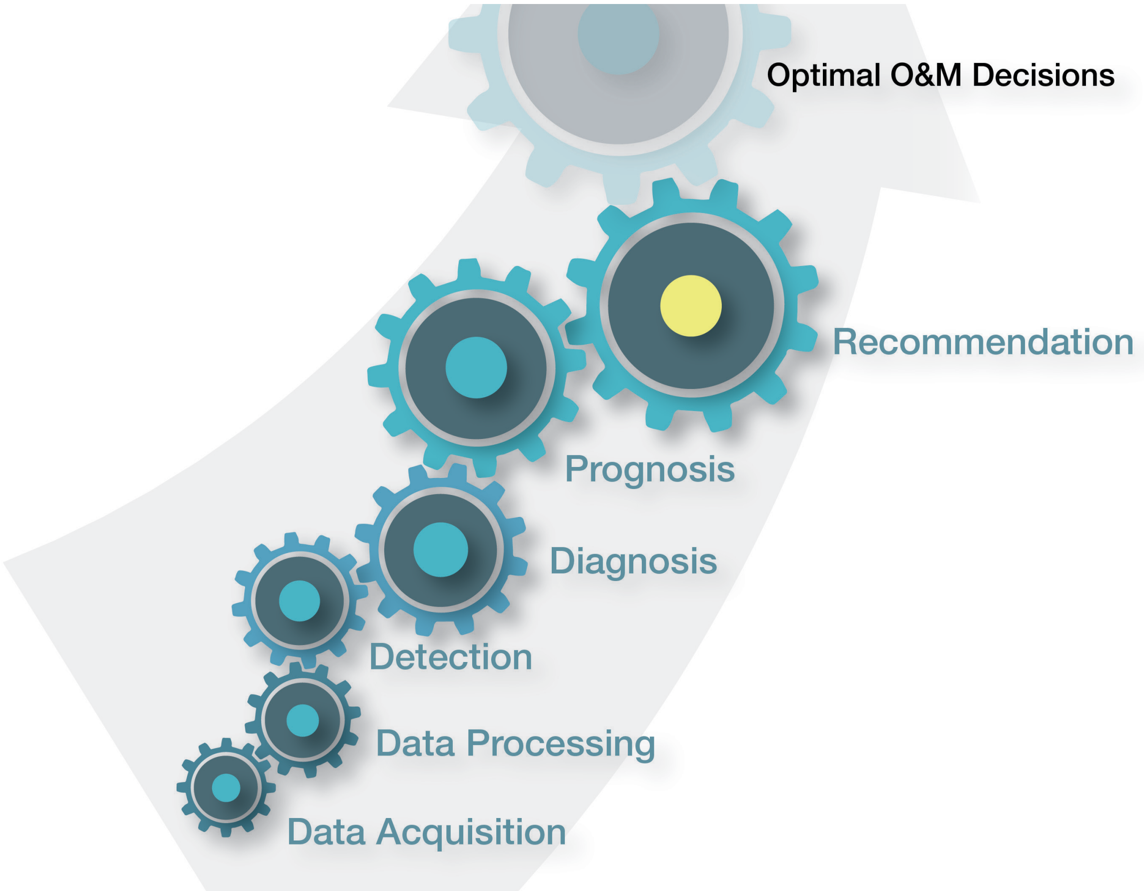


*Figure 1. The six steps of an effective condition monitoring process*

The process moves from "Data Acquisition" to "Recommendation" in a linear fashion. Since the output of each step drives the next, the quality of work done in one affects downstream performance. Thus, considering the CMS from a holistic systems level is most appropriate. It is also worth noting that the condition monitoring process stops at recommendation. It is up to the operator to use those recommendations to make better operations and maintenance decisions.

There are many types of wind turbine CMS with very different methods of providing turbine health information. To illustrate, we will use an example that everyone can recognize: a visit with your physician.

## Data Acquisition

**Every CMS starts with a sensor that** ***translates a physical phenomenon into an analog measurement, which is then converted into digital format*** **for further processing.**

In our example, a doctor takes a blood pressure measurement as a routine part of the visit. The data acquisition sensors, in this case, are the stethoscope used to measure cardiac cycles and the pressure cuff used to measure arterial pressure. The digitization of the sensor outputs is performed by the ears (stethoscope) and eyes (pressure cuff) of the doctor. These are not high-precision sensors, unlike wind turbine CMS—where the fidelity of measurements affects downstream processing.

With a CMS, understanding the sensitivity, bandwidth, and accuracy of the chosen sensor is important, as well as knowing if it can determine component health on all of the fault modes that can affect it. Oil debris systems can detect pitting failures but cannot detect cracking faults. Vibration-based systems can detect both pitting and cracking, but most cannot determine the health of components in the planetary section. Prospective buyers should take an inventory of the components on their wind turbines that have been driving the largest maintenance costs and determine their most common fault modes.

## Data Processing

**After the sensor measurement has been converted to a digital format, the CMS must *process the sensor measurements into meaningful indications of component health*.**

n our example, the physician combines the data from the cardiac cycle (stethoscope output) with the pressure variations (pressure cuff output) measured at the same time. They then average all the pressure variations over the course of the measurement. The complexity of the two measurement signals is reduced to two simple numbers—120/80 mmHg, for example—which characterizes the patient's current blood pressure.

With CMS, the data processing step involves two distinct sub-steps. The first is to *isolate the relevant portion of the measurement signal* from the "noise," which involves some sort of filtering of the original signal. In our example, the physician only looks at the pressure levels during specific parts of the cardiac cycle, filtering out the rest of the extraneous values. When the signal isolation is done well, it increases the sensitivity of the CMS, allowing for easier discrimination between "unfaulted" and "faulted" components. It also reduces the inevitable variation in these component condition indicators due to the complex environments and varied conditions in which wind turbines operate.

The second step of data processing is *extracting the salient features of the filtered signal* that provide an indication of component condition. The resultant indicators should ideally identify the presence of different fault modes in the components. For example, a gear can have several fault modes including root cracks, surface pitting, or misalignment. Each of these fault modes manifest themselves in different ways, so no single condition indicator will accurately characterize all of these faults. Therefore, several indicators based on different filtering methods should be used to identify potential fault modes in each component.

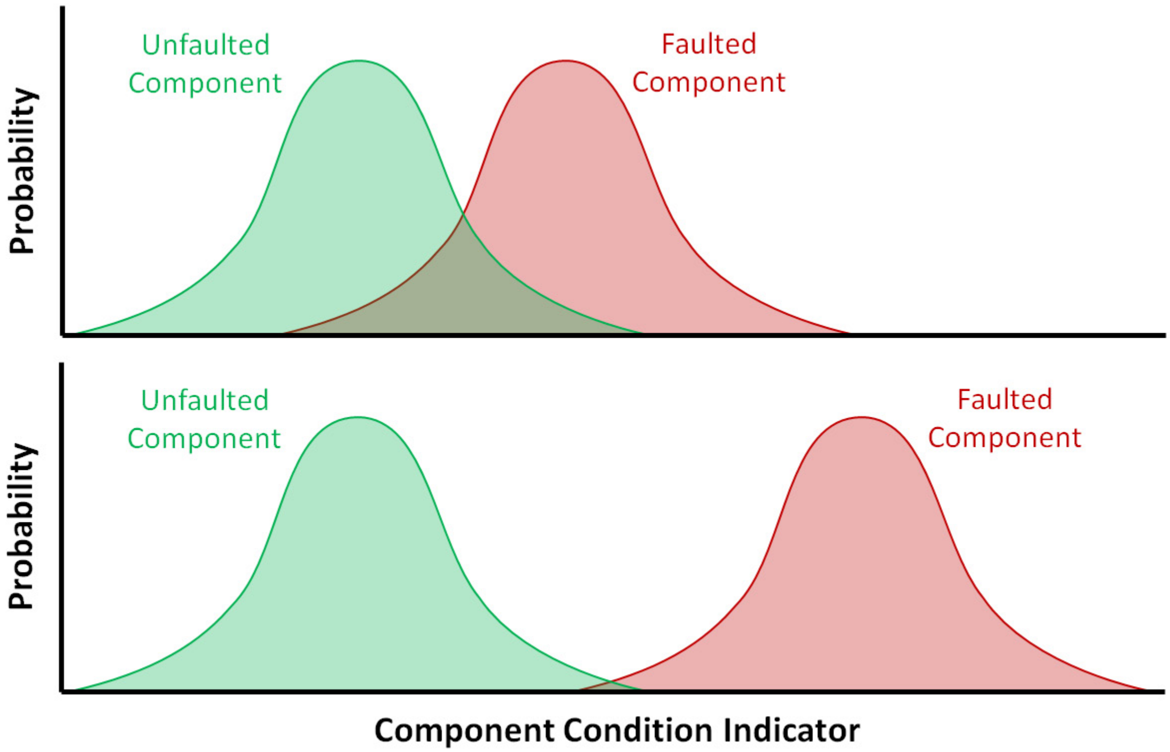


*Figure 2. Effective data processing is the difference between poor fault discrimination (top) and good fault discrimination (bottom)*

When data is processed correctly, the condition indicators identify all potential fault modes and easily discriminate between faulted and unfaulted components.

Figure 2 shows two different scenarios: a condition indicator that results from poor data processing (top), and a condition indicator that results from effective data processing. In both of the graphs, the green distribution is the range of the condition indicator typical for an unfaulted component, while the red distribution is the range of the condition indicator typical for a faulted component.

In the top graph, there is a great deal of overlap between the two distributions due to inadequate data processing, so the ability to discriminate between a faulted and unfaulted component is poor. In the bottom graph, effective data processing has provided adequate separation of the faulted and unfaulted conditions, so discriminating between the two is straightforward.

## Detection

**Once the measured signals have been turned into condition indicators, the CMS must *classify whether the condition indicators are "normal" or "abnormal."***

In our example, the physician has determined the patient's blood pressure, but that measurement itself is not instructive. It is not until that measurement is compared to the commonly used 120/80 mm Hg threshold that we can determine if it is high or not.

With CMS, this is done by comparing the current condition indicator to a reference range, which can either be a statistical baseline or model-based. Setting the level of the threshold used to classify the condition indicators as either "normal" or "abnormal" is the crux of the detection step. The threshold is typically a high limit set on a condition indicator.

In our example, the blood pressure threshold is based on studies including large populations of patients with no known hypertension. For wind turbine condition monitoring, it is much more difficult due to the complexity of the systems and the number of different turbine makes and models in the field.

Figure 3 shows the unfaulted (green curve) and faulted (red curve) component distributions as before, but a fault threshold has been added. In this case, the unfaulted and faulted distributions have significant overlap, so misclassifications are inevitable.

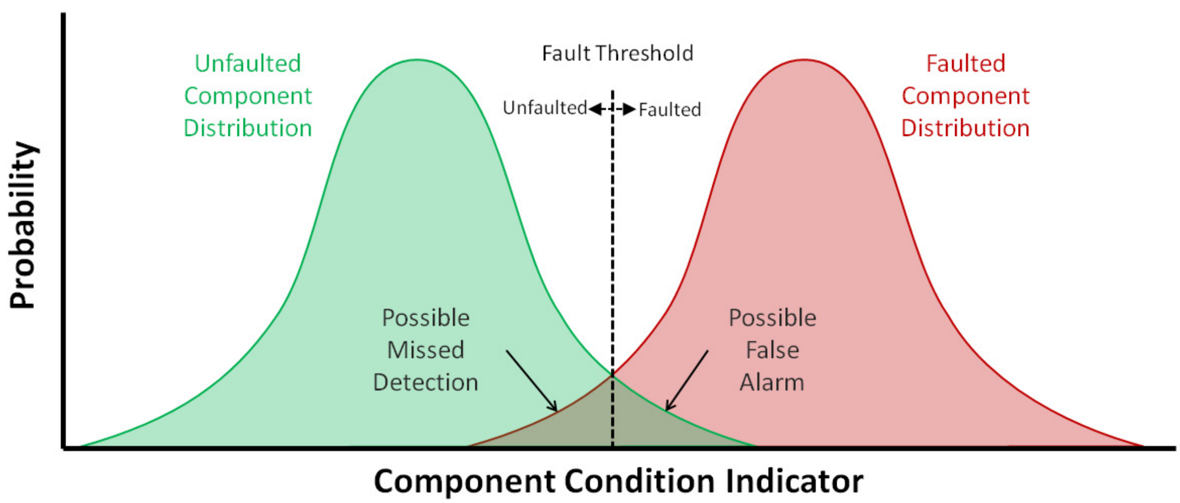


*Figure 3. The threshold level and how it is chosen is critical to system performance.*

In the figure, the threshold was set to balance possible missed detections and false alarms. In practice, setting thresholds is even more difficult because there are few, if any, measurements of what a faulted component looks like. In the best case, CMS thresholds are set based on knowing what an unfaulted component looks like (green distribution in the graph) and a predefined probability of false alarms. Because of the inherent complexity and direct impact threshold setting has on performance, understanding how a condition monitoring provider will set thresholds is one of the critical questions to ask when selecting a system. Systems that use a poor process for setting thresholds are more prone to false alarms, driving unnecessary maintenance.

## Diagnosis

**Once a condition indicator detects a faulted component, the CMS must *validate the fault and determine its location and severity*.**

If a patient had a high blood pressure reading, a physician may be inclined to diagnose hypertension. Yet upon further discussion, it was discovered that the patient had a stressful week. In addition, the patient has no family history of hypertension and their historic blood pressure values are lower. In this case, a physician would not diagnose hypertension, even though high blood pressure was detected. Now imagine if the patient did have a family history of hypertension and their historic blood pressure values had been trending upward for several years. The diagnosis would be hypertension and the next step would be to determine the severity.

This is done by examining the context in which the indication was high. The CMS can compare the current condition indicator to the historic value of the same condition indicator and under what operating condition it occurred. If this is the first high value and it happened under high transient loading, it may be best to ignore this indication until more evidence is gathered.

Just like in the physician example, determining the severity of a turbine fault is a critical part of a CMS diagnosis. The action that is needed for a component with a small fault and months of remaining useful life left is very different then the action needed for a severely faulted component with only hours left.

## Prognosis

**Once the fault has been validated and the severity is known, the next piece of information needed is *an estimate of how much longer that component will last before it needs to be replaced*, also known as the remaining useful life.**

The remaining useful life of the component can be estimated in several ways, but it requires knowledge of two things: what the current severity of the fault is and an estimate of the future operating conditions of the component.

In our example, once the severity of the patient's condition is understood, the physician can determine how it will degrade. If the current hypertension is low (fault severity) and the patient already lives a healthy lifestyle (future operating conditions), the prognosis may be that the hypertension will have little impact on the patient's future well-being. If the hypertension is currently low but the patient lives a sedentary, unhealthy lifestyle, the prognosis may be that if the hypertension is left untreated it will lead to heart disease in five years. Both situations started at the same severity level, but the anticipated future conditions led to vastly different prognoses.

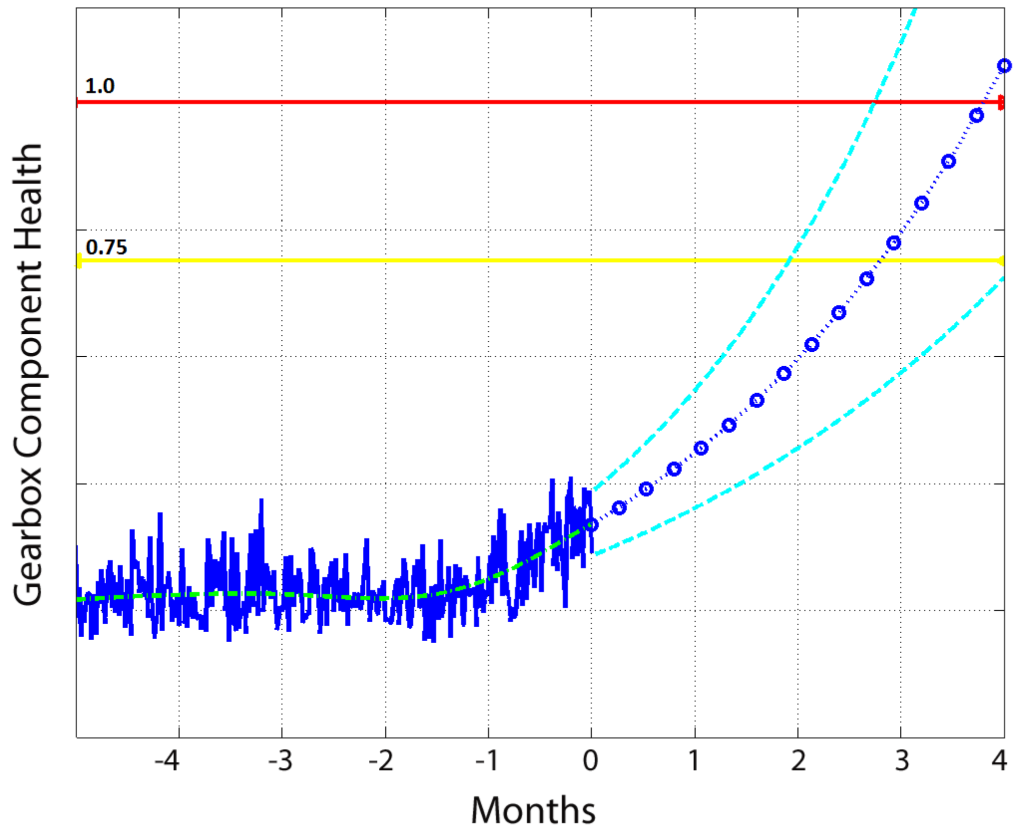


*Figure 4. An estimate of remaining useful life for a component is created by projecting the current health forward using a model of component health and the operating condition it is most likely to see*

A prognosis for a wind turbine CMS is used slightly differently. Instead of changing future operating conditions to prevent component failure, an estimate is used to determine when the component will reach the end of its useful life. Figure 4 show a graphical example of the projection of future component health. The amount of time between the present and the time when the estimated trajectory of the component's future health (blue dotted line) crosses a pre-defined threshold (red line) defines the remaining useful life of the component.

## Recommendation

**Once the CMS has an estimate of fault severity and the remaining useful life of the component,** ***the necessary maintenance action and when it should be performed*** **can now be determined.**

The recommendation step is really an aggregation step; information is taken from the diagnosis and prognosis steps and combined into a clear recommendation of what to do next. In our example, if the patient is diagnosed with mild hypertension and the prognosis is that there will be no impact on their overall health, the recommendation would be to maintain their current lifestyle. If the patient receives the same diagnosis but the prognosis is that the mild hypertension will lead to heart disease in five years, the recommendation would be to exercise more and change to a healthier diet.

For wind turbines, the recommendation comes in the form of a maintenance action that will be required. If a bearing is faulted, the recommendation could be to verify the fault through visual inspection within the next month and schedule a replacement of the bearing within three months. This allows the operator to plan maintenance outages ahead of time, reducing downtime and lost revenue.

## Closing Thoughts

The framework of the condition monitoring process presented here provides a guideline for prospective buyers considering a CMS purchase. Many systems available do not cover the entire condition monitoring process. This may require an operator to interpret a significant amount of data, so be sure to ask vendors what parts of the process their system covers and if additional services are required to get to a recommendation. In the end, the efficacy of a CMS is only as good as its ability to provide operators with information that drives better operations and maintenance decisions.

*About the author:*

Brogan Morton is the Product Manager for TurbinePhD, a vibration-based predictive health monitoring system, at Renewable NRG Systems. Brogan started his career performing engineering research on the diagnostics and prognostics of mechanical components in aerospace systems. Morton holds a Master's degree in mechanical engineering from the University of New Hampshire and an MBA from Idaho State University.