Rethinking Reliability in Terms of Margins

D. Mandelli[1], C. Wang[1], K. A. Manjunatha[2], V. Agarwal[1], L. Lin[1]

[1]*Idaho National Laboratory (INL), Idaho Falls, Idaho, 83415, United States*
Diego.Mandelli@inl.gov, Vivkek.Agarwal@ink.gov, Congjian.Wang@inl.gov, Linyu.Lin@inl.gov

[2]*Amazon Lab126, Sunnyvale, California, 94089, United States*
koushik@crimson.ua.edu

ABSTRACT

Ideally, reliability methods should support assessing and managing system health by utilizing the integrated health information of all the system assets. An important aspect is that reliability data employed in these methods are an approximated integral representation of past industrywide operational experience. Thus, they neglect an asset's present health status (obtainable, for example, from online monitoring data and diagnostic assessments) and forecasted health projection (when available from prognostic models). Asset health should be informed solely by that specific asset's current and past performance data and should not be an approximated integral representation of past industrywide operational experience. Sensor data, diagnostic assessments, and prognostic assessments are in fact not considered in plant reliability models. In addition, propagating quantitative health data from the asset level to the system level is made challenging by the diverse nature and structure of health data elements (e.g., vibration spectra, temperature readings, and expected failure time). Ideally, in a predictive maintenance context, system reliability models would support decision-making by propagating available health information from the asset level to the system level to provide a quantitative snapshot of system health and identify the most critical assets. This paper directly addresses the limitations of current reliability methods by proposing a different approach to reliability modeling: a method that relies on asset diagnostic, prognostic, and monitoring data to measure asset health. Propagating health data from the asset level to the system level is performed through reliability models, not in terms of probability but rather in terms of margin, with margin being the "distance" between the asset's present status and an undesired event (e.g., failure or unacceptable performance).

1. INTRODUCTION

Current reliability approaches assess and quantify the reliability associated with complex systems, such as nuclear power plants (NPPs). These approaches are generally based on classical Boolean logic structures, such as event trees (ETs) and fault trees (FTs) (Rausand, 2020). The outcome obtained by combining FTs and ETs is the set of minimal cuts sets (MCSs), with each MCS representing a unique combination of basic events (BEs) that leads to an undesired outcome (e.g., core damage). The probabilistic evaluation of an MCS is performed by evaluating the product of the probability values associated with each BE. A relevant factor here is that the probability values associated with the BEs used in the plant models are updated at least every 4 years based on past operational experience through the a Bayesian statistical process (Siu, 1998). Hence, the probability value of a BE associated with a physical asset (e.g., a centrifugal pump or motor-operated valve) in no way reflects that asset's actual condition and performance.

This fact plays a major role in the application of plant reliability models to support risk-informed decisions. To reduce operation and maintenance costs, existing NPPs are moving from corrective and periodic maintenance to new types of predictive maintenance strategies (Agarwal, 2021). This transition is designed such that maintenance is conducted only when the asset requires it (i.e., prior to undergoing imminent failure). And though these benefits cannot be achieved through actual reliability modeling methods and currently employed reliability data, they can be achieved by employing asset-monitoring sensors, automated data acquisition systems, data analysis methods, and improved decision-making processes. Combined, these resources can provide precise information on asset health, track its degradation trends, and estimate its expected failure time. Based on such information, maintenance operations can be scheduled and performed for each asset as needed. This dynamic context of predictive maintenance operations requires new methods of data analysis, propagating asset

health information from the asset level to the system level, and optimizing plant resources.

This paper provides an alternative reliability approach for a predictive maintenance context in which a direct link is created between equipment reliability (ER) data and decision-making. Rather than thinking of reliability in terms of system and asset probability of failure, we propose a reliability mindset based on the concept of margin (Mandelli, 2023). An asset's health is quantified by determining its margin, based on the asset's current and historical monitoring data. The margin values of the monitored asset are then propagated through system reliability models (e.g., FTs or reliability block diagrams) to identify the assets that are more critical to guarantee system operation. We show how a margin-based approach can assess asset health, based solely on current and historic monitoring data (e.g., condition-based, anomaly detection, diagnostic, and prognostic data) (Xingang, 2021). The novelty of this margin-based approach is that it directly addresses the limitations of classical reliability modeling approaches by propagating and integrating health data from the asset to the system level.

Here we focus on directly employing ER data effectively optimize maintenance operations (Pinciroli, 2023). As part of this decision-making process, the assessment of asset current and/or future conditions is required; this knowledge can be produced by employing condition monitoring, anomaly detection, or prognostic systems (Zio, 2022). In particular, we are answering this question: how health data/knowledge can be propagated from the asset to the system level? It is common practice to measure asset health using asset health indices (AHIs) (Hjartarson, 2006); the definition of such indices is typically situation and asset dependent (e.g., an AHI might be defined using a color-coded strategy or can be numerically quantified using arbitrary scales). One of the objectives of this paper is to provide a margin-based definition of AHI that is consistent on several operational contexts. The concept of margin is here borrowed from structural reliability analysis theory where margin is defined as the "distance" between the load and resistance probability distribution functions (Melchers, 2018). Reference (Lewis, 2022) presents a bridge between prognostics and health management (PHM) and probabilistic risk assessment (PRA). Our work conceptually differs from (Lewis, 2022) in two elements. The first one is related to the fact that once an asset is supported by a PHM system, the reliability of that asset loses most of its stochasticity since the monitoring activity of such asset is designed to inform on its conditions that might lead to its failure. The second one is related to the kind of decisions that a PHM and PRA are supporting. (Mandelli, 2023) claims that PHM systems support dynamic decisions where maintenance activities are scheduled only when required, while PRA models support static decisions, such as setting periodic surveillance and maintenance activities.

2. MARGIN MODELING

Reference (Mandelli, 2023) expands the meaning of the word "reliability" to better reflect the needs of system health and asset management decision-making processes. Rather than focusing on the likelihood of a given event (in probabilistic terms), we think in terms of how far this event is from occurring. This new interpretation of reliability shifts the focus away from probability of occurrence and toward an assessment of how close an asset is to reaching an unacceptable level of performance or failing (see Fig. 1). Note that two data elements are required for this assessment: the estimated actual health condition of the asset, which can be acquired by the asset-monitoring system or through diagnostic methods, and the limiting conditions that must be avoided, which can be acquired from past operational experience (e.g., monitoring data generated by similar assets under failure conditions).

An asset's margin value M is defined over the $[0,1]$ interval, where $M = 1$ corresponds to a perfectly healthy asset (requiring minimal to no maintenance attention) and $M = 0$ corresponds to a faulty asset (requiring maintenance attention). Figure 1 provides a glimpse (in graphical form) of how a link between monitoring data and decision-making can be established through a margin-based reliability mindset.

Note that margin quantification is impacted by the availability of monitoring data and can be defined over heterogenous variables, such as pressure, vibration spectra, and time. For example, when dealing with condition-based monitoring data (both current and archived), margin M is defined here as the distance between actual and past conditions (e.g., oil temperature and vibration spectrum) that lead to failure (see Fig. 2). Hence, margin-based reliability modeling provides a unified approach to dealing with heterogeneous monitoring data elements.
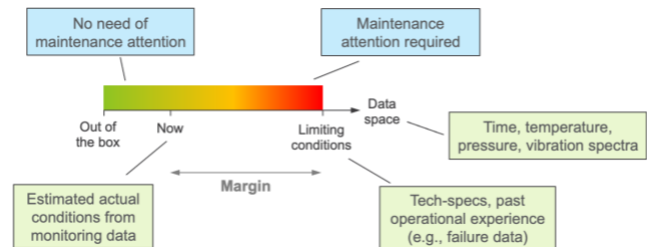

Figure 1. Graphical representation of margin, based on actual asset-monitoring data.

Note that the margin value of an asset is not static but changes with time, depending on asset conditions. For example, if degradation due to usage is observed from the monitoring data, the corresponding asset margin value decreases. Conversely, if a maintenance operation is performed on that same asset (e.g., restoration of centrifugal pump bearings), the asset margin value increases.

This mindset shift regarding the concept of reliability (i.e., margin based instead of probability based) offers the

advantage of directly linking the asset health evaluation process with standard plant processes for managing plant performance (e.g., plant maintenance operations and budgeting processes). The transformation also supports decision-making in a form that is more familiar and readily understandable to plant system engineers and decision makers.

So far, margin has been defined for one single asset; the next step is to quantify the system's margin value after obtaining the margin values of its assets. The propagation of margin values from the asset level to the system level is performed through classical reliability models, such as FTs or reliability block diagrams (Lee, 2011), which are solved using different rule sets (Mandelli, 2023) instead of set theory-based operations.

In this respect, margin-based operators for assets in both series (OR operator) and parallel (AND operator) configurations must be defined. As an example, consider two assets (A and B). The margin M of both assets can be visualized in a 2D space, as shown in Fig. 2. Starting with brand-new assets (i.e., $M_A, M_B = 1$) the aging and degradation that affects both is represented by the blue line, which parametrically signifies the combination of both margins $M_A(t)$ and $M_B(t)$ at a specific point in time t. Note that if no maintenance (preventive or corrective) was ever performed on either asset, this path would move from coordinates $(1,1)$ to coordinates $(0,0)$, where both assets would be considered failed. Hence, the coordinates $(0,0)$ in Fig. 2 represent the event "A AND B." Similarly, when the blue line reaches the x or y axis of Figure 2 (characterized by $M_B = 0$ and $M_A = 0$, respectively), either asset A or B has failed. Hence, the points in Fig. 2 characterized by either $M_B = 0$ or $M_A = 0$ represent the event "A OR B."

Now we can calculate the margin M for the AND and OR events described above. This is accomplished by following the definition of margin: by measuring the distance between the actual condition of assets A and B and the conditions identified by the event under consideration (e.g., the occurrence of both or either event). The margin for A AND B can be calculated as the distance between the current point of coordinates (M_A, M_B) to the point $(0,0)$. The margin for A OR B is the minimum distance from the current point of coordinates (M_A, M_B) to the x or y axis of Fig. 2 (where $M_B = 0$ and $M_A = 0$, respectively):

$$M(A\ AND\ B) = dist[(M_A, M_B), (0,0)] \qquad (1)$$
$$M(A\ OR\ B) = min(M_A, M_B) \qquad (2)$$

where the function $dist[.,.]$ indicates the metric designed for calculating the distance between two points in an Euclidean space (e.g., if Euclidean distance is employed, $M(A\ AND\ B) = \sqrt{M_A^2 + M_B^2}$). Mandelli (2023) provides a set of considerations regarding the choice of appropriate distance metric $dist[.,.]$ to be employed. In summary, Euclidean and Manhattan distance metrics represent the lower and upper bounds for $M(A\ AND\ B)$ (i.e., $\sqrt{M_A^2 + M_B^2} \leq M(A\ AND\ B) \leq M_A + M_B$). If the temporal

evolution of M_A and M_B is available, a more precise estimate of $M(A\ AND\ B)$ can be obtained.
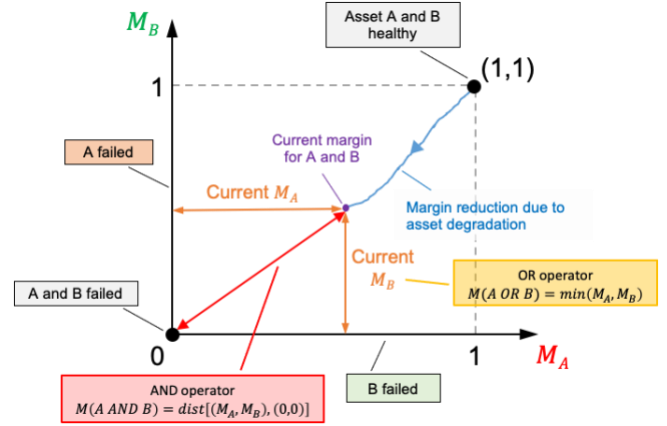

Figure 2. Graphical representation of event occurrences, based on a margin framework.

Eqs. (1) and (2) allow us to propagate margin values through classical reliability models (e.g., FTs or reliability block diagrams) to quantify the system margin M_{sys}. The next step is to determine each asset's importance (in a margin-based reliability context). In a classical reliability setting, this is done by relying on risk-importance measures (Lee, 2011), such as the Birnbaum or Fussell-Vesely measures. Given the different nature of the margin concept, we require a reliability importance measure, here indicated as RIM_α, that captures the impact of asset margin M_α on system margin M_{sys}. Here, we rely on a classical sensitivity measure (derivative based) for an asset α, defined as:

$$RIM_\alpha = \frac{\partial M_{sys}}{\partial M_\alpha} \qquad (3)$$

Simply stated, RIM_α indicates how a small variation of M_α (e.g., improving the health of asset α) directly affects system margin M_{sys}.

3. INTEGRATION OF ER DATA INTO MARGIN MODELS

The definition of margin presented in Section 2 is abstract; an application within a more practical setting depends on the phenomena of interest—and especially the monitoring data available. This section provides more quantitative details on how margin can be quantified depending on the available ER data.

3.1. Technical Specifications Data

As indicated in Section 2, a margin value can be calculated as the distance between the actual and limiting conditions. In practical settings, limiting conditions can be represented by the technical specifications of the considered asset, which are normally provided by the manufacturer. As an example, to ensure the proper function of induction motors, oil viscosity must be below a specified limiting condition. Oil viscosity can significantly change as a function

of motor rotation speed. In this context, asset margin can be calculated as the difference between the limiting condition specified in the technical specifications and the currently measured oil viscosity.

In general terms, given an upper limiting condition x_{LC} for a monitored variable x_{obs}, a margin M can be defined as:

$$M(x_{obs}) = \frac{x_{LC} - x_{obs}}{x_{LC} - \min(x_{obs})} \quad (4)$$

where $\min(x_{obs})$ indicates the minimum allowable value for x_{obs}.

As an example, induction motors are designed to operate within specified differential temperature limits. These limits indicate the maximum permissible difference between the motor temperature and environmental temperatures that various classes of insulation materials can withstand (this temperature limit can range from 80°C to 120°C, depending on the insulation material). In this scenario, x_{LC} is represented by the specified temperature limit, while x_{obs} is the difference between the actual motor temperature and environmental temperature.

3.2. Observed Reliability Parameters

Current industrywide available datasets often report the mean time to failure (MTTF) values for assets, given the past operational experience of similar assets operating under similar environmental conditions (e.g., temperature and humidity). In this context, no monitoring data are available, and only past operational experience can be used. Similar to the reasoning behind Eq. (4), based on the asset's current age t (since installation or refurbishment) and estimated $MTTF$, its margin can be defined as a linear function of t:

$$M(t) = \begin{cases} \frac{MTTF - t}{MTTF} & if \ t < MTTF \\ 0 & if \ t \geq MTTF \end{cases} \quad (5)$$

When the considered asset is brand new (i.e., $t = 0$), margin $M = 1$. When the same asset is approaching its estimated $MTTF$, margin becomes $M = 0$.

3.3. Condition-based Data: Healthy Data

Here, we consider a case in which the available monitoring data for the asset being considered were collected exclusively when the asset was healthy $\Xi^{healthy}$, meaning that data pertaining to asset degradation or failure are unavailable. $\Xi^{healthy}$ represents a collection of past observation data elements ξ^{obs}. The following notation is used throughout this paper: a single observation data element ξ^{obs} can be composed of L observed variables x_l ($l = 1, \dots, L$) (i.e., $\xi^{obs} = [x_1, \dots, x_L]$), and the nature of the observed variables x_l can be heterogenous in nature (e.g., temperature, pressure).

In this kind of situation, an asset's health status can be established by measuring how actual monitoring data differ (distance-wise) from healthy data. In this respect, anomaly detection tools (Nassif, 2021) designed to quantify the residual between the actual observed data ξ^{obs} and the

predicted data ξ^{rec} (which are computed from ξ^{obs} and $\Xi^{obs-healthy}$) can be employed. Such tools can be based on a kernel density estimation, for example the auto-associative kernel regression method (Baraldi, 2015) or on deep-learning-based methods, e.g., see (Zhang, 2019). Under normal conditions, ξ^{rec} is very similar to ξ^{obs} (i.e., $\xi^{obs} \cong \xi^{rec}$). $\xi^{obs} \neq \xi^{rec}$ indicates anomalous behavior (e.g., asset degradation).

In this context, a margin value can then be defined by measuring the difference between ξ^{rec} and ξ^{obs} as:

$$M(\xi^{obs}) = e^{-\left(\frac{\|\xi^{obs} - \xi^{rec}\|}{h}\right)^2} \quad (6)$$

where $\|\xi^{obs} - \xi^{rec}\|$ indicates the residual between the observed and predicted data and h represents the comparison parameter between ξ^{rec} and ξ^{obs} (expressed in terms of standard deviation). When the asset is experiencing normal conditions, $\xi^{obs} \cong \xi^{rec}$, $M = 1$. If the asset is experiencing abnormal conditions, the norm of the difference between ξ^{obs} and ξ^{rec} increases; consequently, M drops to 0.

Note that $\Xi^{healthy}$ is here assumed to cover all possible healthy asset conditions. If this is not the case, when ξ^{obs} enters an unforeseen healthy condition, the obtained margin value will show the asset to be unhealthy. However, once newly observed healthy conditions are recorded, they can be added to the original dataset $\Xi^{obs-healthy}$.

An example is shown in Fig. 3, which reflects a set $\Xi^{healthy}$ of observed data elements $\xi^{obs} = [x_1, x_2]$ being collected (the green dots in the top image of Fig. 3). Actual observed data ξ^{obs} are constantly recorded, while ξ^{rec} are determined based on ξ^{obs} and $\Xi^{healthy}$ (see the black and red lines in the top image of Fig. 3), using the auto-associative kernel regression method (Baraldi, 2015). Applying Eq. (6) to this test case makes it possible to generate a temporal profile for the corresponding margin (bottom plot of Fig. 3).

3.4. Condition-based Data: Healthy and Faulty Data

This case extends the one described in Section 3.3 (in which only data generated under healthy conditions $\Xi^{healthy}$ are available) by incorporating data generated under faulty conditions, indicated here as Ξ^{faulty}. It is assumed that, in the presence of an asset fault, the actual observed data ξ^{obs} can be seen transitioning from $\Xi^{healthy}$ to Ξ^{faulty}.

In this scenario, by following the definition of margin given in Section 2 and by being provided with actual observed data (containing both historic healthy $\Xi^{healthy}$ and faulty data Ξ^{faulty}), a margin value can be determined by comparing the mutual distance of ξ^{obs} from the two populations: $\Xi^{healthy}$ and Ξ^{faulty} (see Fig. 4). In mathematical form, a margin can be written as:

$$M(\xi^{obs}) = \frac{D(\xi^{obs}; \Xi^{faulty})}{D(\xi^{obs}; \Xi^{faulty}) + D(\xi^{obs}; \Xi^{healthy})} \quad (7)$$

where the operator $D(.;.)$ represents the distance one single data element (i.e., ξ^{obs}) and a population of data elements

(either $\Xi^{healthy}$ or Ξ^{faulty}). The choice of operator $D(.;.)$ may depend on several factors, as dictated by the distribution of the $\Xi^{healthy}$ and Ξ^{faulty} populations in the data space.
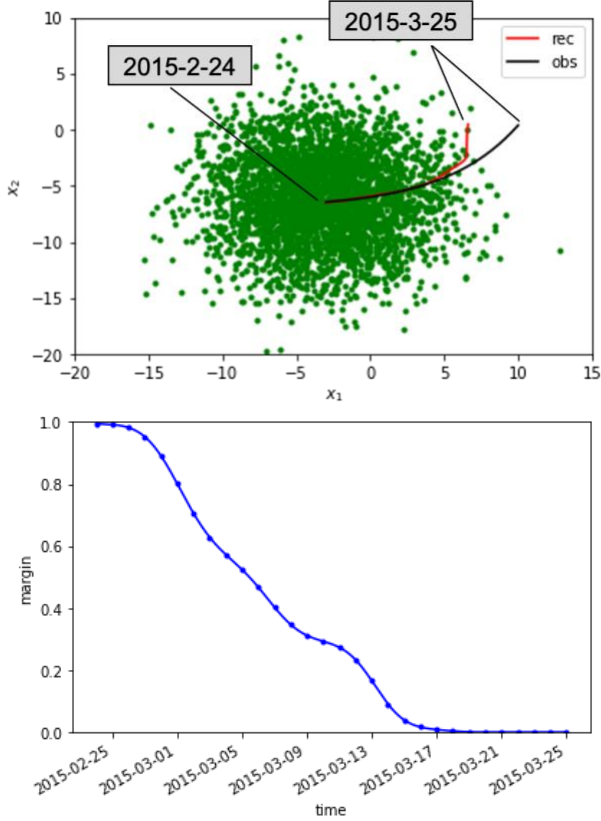

Figure 3. (Top) Representation of $\Xi^{healthy}$ (green population), ξ^{obs} (black line), and ξ^{rec} (red line) in the x_1, x_2 space. (Bottom) Temporal profile of the corresponding margin.
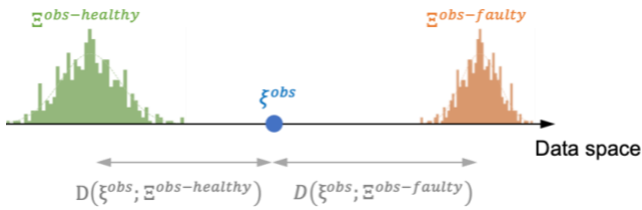

Figure 4. Margin calculation, given the current status of the monitored asset ξ^{obs} when both healthy $\Xi^{healthy}$ and faulty Ξ^{faulty} data are available in the $[x_1, .., x_L]$ data space.

Note that a distance-based approach for $D(\xi^{obs}; \Xi)$ is effective when the healthy and faulty data are well separated from each other in the $[x_1, .., x_L]$ space. In practical scenarios, however, these two populations of data elements may overlap. In such cases, margin can be quantified by using density-based methods (Hastie, Tibshirani, and Friedman, 2001), which are designed to translate (e.g., via kernel density estimation methods) the $\Xi^{healthy}$ and Ξ^{faulty} datasets into

probability distribution functions (PDFs): $pdf^{healthy}$ and pdf^{faulty}. Then, given a current observed measurement ξ^{obs}, margin can be quantified by evaluating these two PDFs at the coordinate ξ^{obs}:

$$M(\xi^{obs}) = \frac{pdf^{healthy}(\xi^{obs})}{pdf^{healthy}(\xi^{obs}) + pdf^{faulty}(\xi^{obs})} \qquad (8)$$

This equation weighs the PDF values at coordinate ξ^{obs} for both the healthy and faulty conditions. When ξ^{obs} is located in a region of the $[x_1, .., x_L]$ space dominated by healthy data, $pdf^{healthy}(\xi^{obs}) \gg pdf^{faulty}(\xi^{obs})$, and $M(\xi^{obs}) \cong 1.0$. Conversely, when ξ^{obs} is located in a region of the $[x_1, .., x_L]$ space dominated by faulty data, $pdf^{healthy}(\xi^{obs}) \ll pdf^{faulty}(\xi^{obs})$ and $M(\xi^{obs}) \cong 0.0$.

Figure 5 illustrates an example that extends the one shown in Section 3.3. In Fig. 5, $\Xi^{healthy}$ and $\Xi^{ofaulty}$ are shown in the top plot, ξ^{obs} is represented as the black line moving from left to right, and the corresponding margin is shown in the bottom plot. Here, $pdf^{healthy}(\xi^{obs})$ and $pdf^{faulty}(\xi^{obs})$ were generated using kernel density estimation methods (Hastie, Tibshirani, and Friedman, 2001).

An alternative formulation to Eq. (8) can be derived when machine learning (ML) methods (Mohri, 2012) are employed. In this setting, a supervised ML model (i.e., a classifier) is trained using both the faulty and healthy datasets ($\Xi^{faulty}, \Xi^{healthy}$) and is employed to predict, given ξ^{obs}, the class out (either faulty or healthy) to which ξ^{obs} belongs. Such a prediction can be augmented by also determining the probability estimate $Prob^{detec}$ associated with the prediction out. If the [0,1] margin interval is divided into two equally long segments, we can assign the "healthy" class to the [.5,1] interval and the "faulty" class to the [0,.5] interval. Hence, the predicted class out generated by the ML model determines the margin variability interval (either $[0,.5]$ or $[.5,1]$). The variable $Prob^{detec}$ (see Fig. 6) is essentially a measure of the prediction accuracy. More precisely, a high value of $Prob^{detec}$ implies a high degree of accuracy in the prediction; conversely, a very low value implies low accuracy. In this context, $Prob^{detec}$ is used to determine the precise margin location in the $[0,.5]$ or $[.5,1]$ intervals. A high value of $Prob^{detec}$ would drive the margin toward the extremes of the intervals (either 0 or 1), whereas a low value of $Prob^{detec}$ would drive the margin toward the common point of the intervals (i.e., 0.5).

Consequently, provided ξ^{obs} and a ML model that can generate both out and $Prob^{detec}$, a margin value can be defined as:

$$M(\xi^{obs}) = \begin{cases} 0.5 - \frac{Prob^{detec}}{2} & if\ out = faulty \\ 0.5 + \frac{Prob^{detec}}{2} & if\ out = healthy \end{cases} \qquad (9)$$

Deep-neural-network-based models (Hastie, Tibshirani, and Friedman, 2001) are an ML model class widely employed for diagnostic applications. Given ξ^{obs}, this class of classifier models generates the class out (either faulty or healthy) to which ξ^{obs} belongs, along with a probability value associated

with each class: $Prob^{healthy}$ and $Prob^{faulty}$ (rather than a single probability value $Prob^{detec}$). Note that, if two classes are considered (faulty and healthy), $Prob^{healthy} + Prob^{faulty} = 1$. The variable out is determined as:

$$out = \begin{cases} healthy & if\ Prob^{healthy} > Prob^{faulty} \\ faulty & if\ Prob^{healthy} < Prob^{faulty} \end{cases} \quad (10)$$
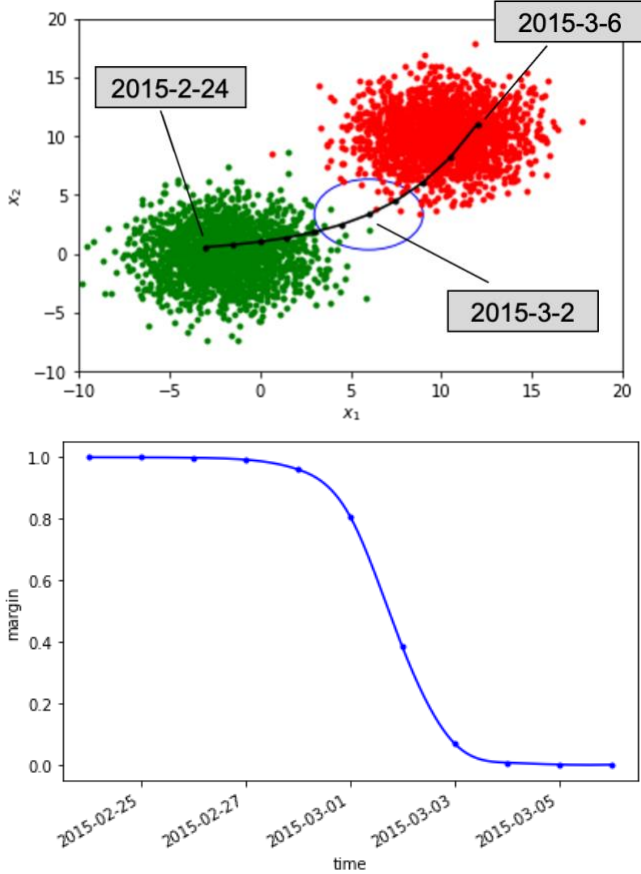

Figure 5. (Top) Representation of $\Xi^{healthy}$ (green population), Ξ^{faulty} (red population), and ξ^{obs} (black line) in the x_1, x_2 space. (bottom) Temporal profile of the corresponding margin.
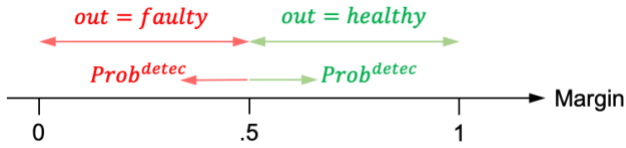

Figure 6. Graphical representation of margin based on the out and $Prob^{detec}$ provided by a ML model.

In this context, margin quantification directly employs the two generated probability values (i.e., $Prob^{healthy}$ and $Prob^{faulty}$) as:

$$M(\xi^{obs}) = Prob^{healthy} = 1 - Prob^{faulty} \quad (11)$$

3.5. Prognostic Data

Estimating an asset's remaining useful life (RUL) provides valuable information regarding the temporal occurrence of the loss of function for the considered asset. Given the stochastic nature of the failure phenomena, RUL is typically expressed in terms of a probabilistic distribution along the temporal axis. Many methods have been developed in the literature to predict RUL for specific assets, and Ferreira and Gonçalves (2022) summarize the most widely used methods. To integrate the RUL PDF (indicated here as PDF_{RUL}) into a margin-based reliability model, we apply reasoning similar to that presented in Section 2. Here, a margin is the distance between the actual time and predicted RUL. The main differences are that the RUL is estimated once a degradation mechanism has been identified (e.g., through an anomaly detection method) and is an actual distribution function rather than a point value.

Once the RUL PDF is estimated, the corresponding margin value can be estimated via two approaches. The first defines the margin as:

$$M(t) = 1 - CDF_{RUL}(t). \quad (12)$$

where CDF_{RUL} indicates the cumulative distribution function corresponding to PDF_{RUL}. The second approach estimates margin as the distance between the actual asset life and a point estimate of the RUL distribution (e.g., the 5th percentile $p_{5\%}^{RUL}$):

$$M(t) = \frac{p_{5\%}^{RUL} - t}{p_{5\%}^{RUL}} \quad (13)$$

where $p_{5\%}^{RUL}$ indicates the 5th percentile of the RUL distribution PDF_{RUL}.

4. TEST CASE: CIRCULATING WATER SYSTEM (CWS) SYSTEM

To develop initial methods and models, a CWS at a Public Service Enterprise Group Nuclear, LLC owned plant site was selected as the target plant asset. The CWS is an important non-safety-related system. As the heat sink for the main steam turbine and associated auxiliaries, the CWS is designed to maximize steam power cycle efficiency (Agarwal, 2021). A CWS consists of the following major equipment (Agarwal, 2021):

- Vertical, motor-driven circulating water pumps (CWPs), each with an associated fixed trash rack and traveling screen at the pump intake to filter out debris and marine life

- Main condenser

- Condenser waterbox air removal system

- Circulating water sampling system

- Screen wash system

- Necessary piping, valves, instrumentation, and controls to support system operation.

The selected plant site (a two-unit pressurized-water reactor) features six circulators at each unit. Schematic representations of the main condensers for Plant Site Unit 2 are shown in Fig. 7.

In this research, the project team focused on optimizing the maintenance strategy for the CWS. To differentiate between motor and pump maintenance activities for each circulator, those assets are hereafter referred to as the CWP motor and the CWP, respectively.

The Unit 1 and Unit 2 CWS process data are collected once per minute and stored in the Plant Site 1 monitoring system. Due to file size restrictions, the project team received CWS process data hourly for both units, ranging from 2009 to 2019. The process data include:

- Ambient air temperature (°F)

- CWP inlet and outlet river temperature (°F)

- CWP motor status (ON or OFF)

- CWP motor stator winding temperature (°F)

- CWP motor inboard-bearing (MIB) and outboard-bearing (MOB) temperature (°F)
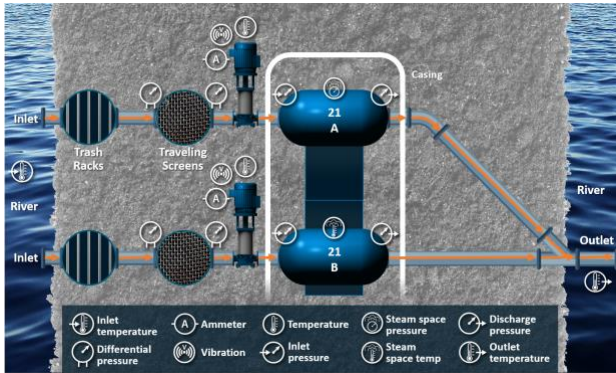
- CWP motor current (amps).

Figure 7. Plant Site Unit 2 CWP combination of 21A and 21B, with sensors and instrumentation.

4.1. Data Processing

As indicated by Agarwal (2021a.), the raw data collected from the NPP are distributed over several data sources and were processed by completing the following steps:

- Text data are converted into numeric form (e.g., the ON/OFF data element is converted into 0/1)

- New features are generated (e.g., pump differential temperature [DT], pump age since refurbishment)

- Pump vibration data are processed through a fast Fourier transform algorithm, and the magnitude of

the vibration signal for specific frequencies is captured

- Based on the system operational history (e.g., maintenance records), data elements are labeled (as pertaining to either healthy or faulty state)

- Missing data entries are resolved

- Data conflicts between the operational history and recorded numerical values are resolved

- All data sources are merged into a single time series

- All features of the time are Z-normalized (each feature x is transformed into $\tilde{x} = \frac{x - mean(x)}{std\_dev(x)}$, where the operators $mean(x)$ and $std\_dev(x)$ correspond to the mean value and standard deviation of the considered variable x, respectively).

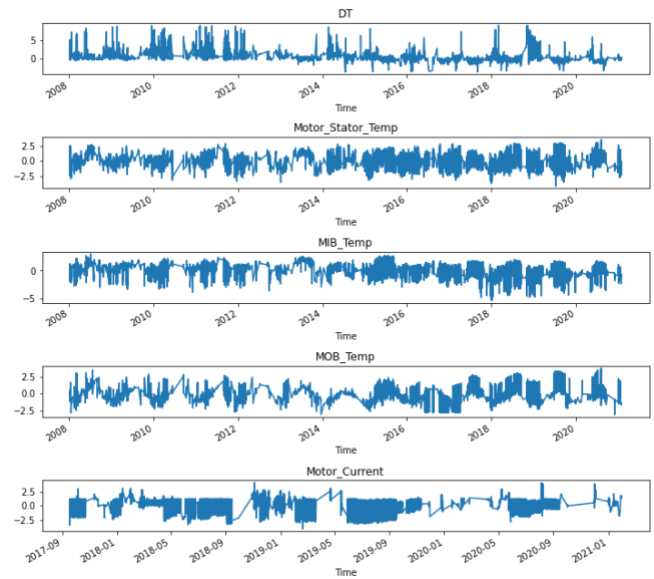A series of preprocessed time series plots is shown in Fig. 8.

Figure 8. Plot of five features of the preprocessed time series. Note that online motor current data are available from 2017, whereas process variables are available from 2009.

For the specific case, by looking at the operational history of the CWS system, we were able to label portions of the data under healthy and faulty conditions (i.e., we were in the scenario described in Section 3.4, if this were not possible, we would have relied only on the known healthy conditions – see Section 3.3). In this respect, Fig. 9 shows box plots of four of the considered features for the healthy and faulty states. These variables were chosen based on their coverage of all healthy and faulty states. Note that the structure of the box plots shifts between healthy and faulty states. This is essential for correctly capturing system health from the available monitoring data. Note that the box plots in

Fig. 9 enable comparison between healthy and faulty states by looking at the distribution of one individual feature at a time.
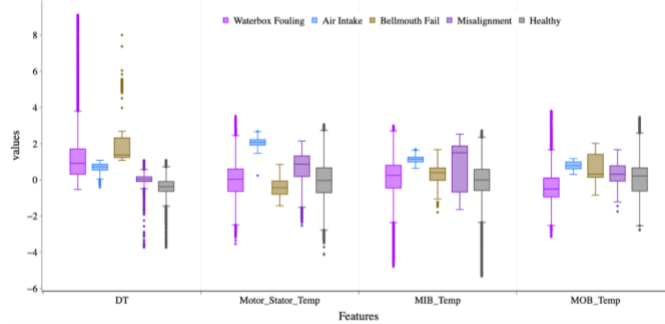

Figure 9. Box plots of four of the considered features (DT, motor stator temperature, MIB temperature, and MOB temperature) for healthy and failure states.

4.2. Margin Model for Air Intake and Misalignment

Given the provided context, both Ξ^{faulty} and $\Xi^{healthy}$ data are available; hence, we employed the density-based method described in Section 3.4 to estimate the $pdf^{healthy}(\xi^{obs})$ and $pdf^{faulty}(\xi^{obs})$. For this specific test case, we considered a subset of the original data points contained in Ξ^{faulty} and $\Xi^{healthy}$ over four monitored variables (i.e., DT, motor stator temperature, MIB temperature, and MOB temperature). We considered a CWS snapshot in which an air intake instance was observed (between May 15 and July 8, 2008). Directly applying Eq. (7) to each ξ^{obs} made it possible to determine the corresponding margin value (see Fig. 10). This plot shows the initial situation, in which the system is in a healthy state ($M^{sys} = 1$), before the margin then rapidly plummets once the faulty condition is initiated. The two peaks that follow were generated during the repair time window.
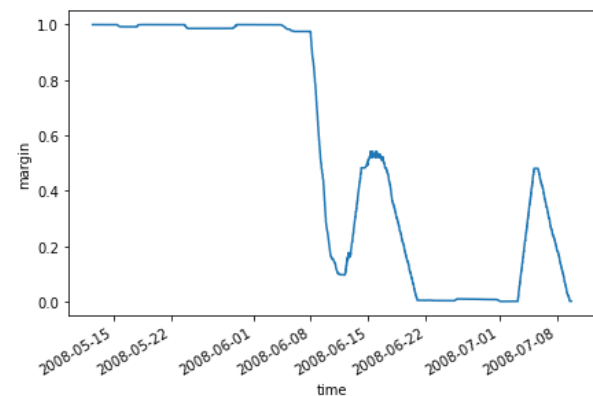

Figure 10. Graphical representation of the margin for air intake during an air intake occurrence, using the density-based method.

An important element to highlight here is that the populations Ξ^{faulty} and $\Xi^{healthy}$ for air intake are fairly

separated, as shown in Fig. 9. This allows us to assign a margin value of $M^{sys} = 1$ when ξ^{obs} is located near the $\Xi^{healthy}$ population, and $M^{sys} = 0$ when ξ^{obs} is located near the Ξ^{faulty} population.

A similar situation can be generated for the misalignment failure mode. In this case, however, the populations Ξ^{faulty} and $\Xi^{healthy}$ are not completely separated as they were shown in Fig. 9. This situation is not uncommon and may be caused by the labeling process applied to the original data (healthy vs. misalignment). We considered a CWS snapshot in which a misalignment instance was observed (between April 2013 and January 2015). Directly applying Eq. (7) to each ξ^{obs} made it possible to determine the corresponding margin value (see Fig. 11). This plot shows that the initial situation, in which the system is in a healthy state, is actually characterized by $M^{sys} = .8$ (instead of $M^{sys} = 1$). This is caused by the fact that the distributions associated with the two populations ($\Xi^{healthy}$ and Ξ^{faulty}) for the misalignment failure mode share some degree of overlap (see also Fig. 9). If the distributions of these two populations do not overlap, $M^{sys} = 1$ when the system is in a healthy state.
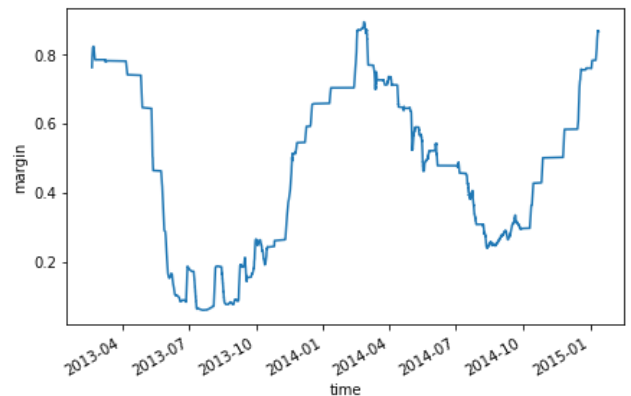

Figure 11. Graphical representation of the margin for misalignment during a misalignment occurrence.

4.3. Margin Model from ML Models

As indicated in (Agarwal, 2021a; 2021b), the following two ML models were generated to perform health and fault classification:

- *Binary classifier*: This module is a XGBoost binary classifier. With CWP data, it predicts whether the CWP is experiencing normal operation or undergoing any degradation at the pump, motor, or system levels. The model is developed by considering time domain features extracted from vibration data, along with the features extracted from monitoring data. Features such as motor current and vibration data are unavailable prior to September 2017 and October 2019, respectively. The missing features are mapped with NaN values. While training and making predictions,

the XGBoost model discards all features with NaN values.

- *Diagnostic model*: This module is a multiclass classifier. For CWP data, it predicts the type of fault a CWP is currently undergoing. The model is developed by considering frequency domain features extracted from vibration data, along with the features extracted from the raw data. Features such as motor current and vibration data are unavailable prior to September 2017 and October 2019, respectively. The missing features are mapped with NaN values. While training and making predictions, the XGBoost model discards all features with NaN values. Figure 12 shows an example prediction generated by the diagnostic model.
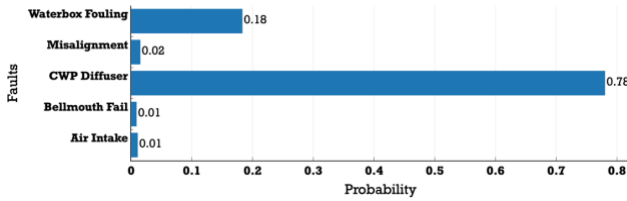

Figure 12. Example prediction by the diagnostic model.

The outputs of these two ML models have been merged to assess the margin for each failure mode by using Eqs. (9)–(11), as indicated in Section 3.4. Figure 13 presents the margin associated with air intake when using ML models; the same temporal profile can be compared against the one shown in Fig. 10, in which a density-based approach was applied to the same dataset. This margin calculation was applied to a subset of observation data ξ^{obs} showing a transition from a healthy state to an air intake faulty state. This transition is captured in a margin sense by observing how the CWS margin for air intake drops from about 0.9 (system healthy) to 0.08 (system in an air intake faulty state).
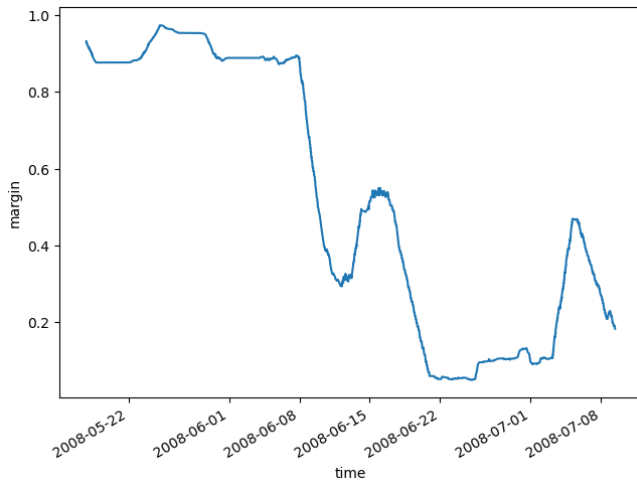

Figure 13. Graphical representation of the margin for air intake during an air intake occurrence, using ML models.

5. CONCLUSION

This paper has described a reliability approach designed to directly employ available condition-based, diagnostic, and prognostic data. It proposed a margin-based approach for assessing asset health, which is based solely on current and historic monitoring data (e.g., condition-based, anomaly detection, diagnostic, and prognostic data). We provided details on how heterogenous ER data elements are employed to assess the status of an asset through a margin value that serves as an analytical measure of its health. We then showed how, depending on the operational context of the asset (e.g., type of failure modes) and the available pertaining to it, a margin value can be quantified using well-known statistical and ML algorithms.

Assessing system health is performed by propagating, through classical reliability models (e.g., FTs or reliability block diagrams), the margin values of those assets that support system function(s). Such propagation is not performed through set theory-based rules but rather through distance-based operations. This information can then be used to assess the reliability importance of each asset in order to identify the most critical assets. A margin-based approach directly addresses the limitations of classical reliability modeling approaches and provides a snapshot of system health—given the availability of monitoring data. These two different approaches are designed to address different types of decisions: classical reliability models support *static* decisions (e.g., a set frequency of periodic maintenance or surveillance operations) based on past operational experience, whereas a margin-based approach directly supports *dynamic* decisions involving maintenance operations that should only be performed when necessary, based on monitoring data (i.e., a predictive maintenance context). Note that the application of these two decision types (*static* and *dynamic*) is dictated by the degradation process being considered. When asset failure occurs suddenly, or the monitoring system cannot capture asset degradation, classical reliability approaches can be used to set preventive maintenance and periodic surveillance frequencies. On the other hand, when assets progressively degrade and the installed monitoring system is able to capture the degradation trend, a predictive maintenance context that relies on a margin-based approach can be set.

An analysis of the CWS system of an existing power plant generated insights into the structure and operational context of real data. Developed statistical and ML methods were employed to assess system health via margin-based operations.

A margin-based interpretation of reliability shifts the focus of the concept away from the probability of occurrence and toward assessing how far away (or close) an asset is to reaching an unacceptable level of performance or undergoing failure. This shift in focus provides a direct link between the asset and system health evaluation process and standard plant processes for managing performance (e.g., plant maintenance

and budgeting processes). It also supports decision-making in a predictive maintenance context in a form that is more familiar and readily understandable to plant system engineers and decision makers (Xingang, 2021).

REFERENCES

Agarwal, V. Araseethota Manjunatha, K., Smith, J. A., Gribok, A. V., Yadav, V., Palas, H., Yarlett, M., Goss, N., Yurkovich, S., Diggans, B., Lybeck, N. J., Pennington, M., & Zwiryk, N. (2021a). Machine learning and economic models to enable risk-informed condition based maintenance of a nuclear plant asset. Idaho National Laboratory Technical Report, INL/EXT-21-61984. https://www.osti.gov/servlets/purl/1770866.

Lewis, A., Groth, K. M. (2022). "Metrics for evaluating the performance of complex engineering system health monitoring models, Reliability Engineering & System Safety, vol. 223. https://doi.org/10.1016/j.ress.2022.108473.

Hjartarson, T., Shawn, O. \ (2006). "Predicting Future Asset Condition Based on Current Health Index and Maintenance Level," in ESMO 2006–2006 IEEE 11th International Conference on Transmission & Distribution Construction, Operation and Live-Line Maintenance.

Mandelli, D., Wang, C., & Hess S. (2023). On the Language of Reliability: A System Engineer Perspective. *Nuclear Technology*. https://doi.org/10.1080/00295450.2022.2143210.

Melchers, R., Beck, A. (2018). *Structural Reliability Analysis and Prediction*. Wiley ed., 3rd edition.

Pinciroli, L., Baraldi, P., Zio, E. (2023). "Maintenance optimization in industry 4.0," Reliability Engineering & System Safety, vol. 234. https://doi.org/10.1016/j.ress.2023.109204.

Rausand, M., Barros, A., & Hoyland, A. (2020). *System Reliability Theory: Models, Statistical Methods, and Applications*. Wiley. https://doi.org/10.1002/9781119373940.

Siu, N., & Kelly D. (1998). Bayesian Parameter Estimation in Probabilistic Risk Assessment. *Reliability Engineering and System Safety*, 62 (1–2), pp. 89-116. https://doi.org/10.1016/S0951-8320(97)00159-2.

Xingang, Z., Kim, J., Warns, K., Wang, X., Ramuhalli, P., Cetiner, S., Kang, H. G., & Golay, M. (2021). Prognostics and Health Management in Nuclear Power Plants: An Updated Method-Centric Review with Special Focus on Data-Driven Methods. *Frontiers in Energy Research*, 9, 696785. https://doi.org/10.3389/fenrg.2021.696785.

Zio, E. (2022). "Prognostics and Health Management (PHM): Where are we and where do we (need to) go in theory and practice," Reliability Engineering & System Safety, vol. 218. https://doi.org/10.1016/j.ress.2021.108119.