

A Method for Automated Cavitation Detection with Adaptive Thresholds

Seth W. Gregg¹, John P.H. Steele², and Douglas L. Van Bossuyt^{3*}

¹ Logical Systems, LLC., Golden, Colorado, 80401, USA
sgregg@logicalsysinc.com

² Department of Mechanical Engineering, Colorado School of Mines, Golden, Colorado, 80401, USA
jsteele@mines.edu

³ Systems Engineering Department, Naval Postgraduate School, Monterey, California, 93943, USA
douglas.vanbossuyt@nps.edu

ABSTRACT

Hydroturbine operators who wish to collect cavitation intensity data to estimate cavitation erosion rates and calculate remaining useful life (RUL) of the turbine runner face several practical challenges related to long term cavitation detection. This paper presents a novel method that addresses these challenges including: a method to create an adaptive cavitation threshold, and automation of the cavitation detection process. These two strategies result in collecting consistent cavitation intensity data. While domain knowledge and manual interpretation are used to choose an appropriate cavitation sensitivity parameter (CSP), the remainder of the process is automated using both supervised and unsupervised learning methods. A case study based on ramp-down data, taken from a production hydroturbine, is presented and validated using independently gathered survey data from the same hydroturbine. Results indicate that this fully automated process for selecting cavitation thresholds and classifying cavitation performs well when compared to manually selected thresholds. This approach provides hydroturbine operators and researchers with a clear and effective way to perform automated, long term, cavitation detection, and assessment.

1. INTRODUCTION

Hydroturbines produce 6.3% of all electrical generation and 48% of renewable energy in the USA (U.S. Energy Information Administration, 2015). While hydro power plants have existed for well over 100 years, issues such as cavitation damage to hydroturbine runners remain problematic for

plant operators (Khurana, Navtej, & Singh, 2012). This paper presents a method to automatically detect damaging cavitation events using existing installed sensors whose data are used to recalibrate the cavitation detection algorithm using hydroturbine ramp-down or ramp-up. This is of particular interest to hydro plant operators since it eliminates required user input and hydroturbine downtime.

The underlying motivation for this work is to reach the goal of estimating remaining useful life (RUL) for hydroturbines, specifically when there is cavitation erosion. It is common practice to use a fixed schedule (based on operating experience) to repair cavitation damage on hydroturbine runners. If RUL can be accurately estimated, then condition-based maintenance of hydroturbines can be implemented. The necessary steps to develop RUL predictions for hydroturbines are as follows:

1. Select a sensor-based cavitation detection method for identifying erosive cavitation and measuring its intensity.
2. Collect cavitation intensity data for a test period that is long enough for accumulative cavitation damage to be measured.
3. Measure the runner material loss over the test period and correlate the loss with the measured cavitation intensity over the same period.
4. Create an erosion rate model for use in estimating runner RUL at any future state based on accumulated cavitation intensity.

It is important to note that a significant amount of data is required including: 1) cavitation detection data, 2) cavitation intensity data, and 3) runner material loss data. These data would then be correlated to develop an erosion rate model to estimate runner RUL. The complexities involved with tracking cavitation detection and intensity data for long periods

*Corresponding Author
Seth W. Gregg et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

in industrial environments have historically been a barrier to creating a prognostic model. For instance, many indicators sensitive to the onset of cavitation, also called a cavitation sensitivity parameter (CSP) as first introduced in (McKee et al., 2015), have specific hardware requirements such as using specialty sensors or high speed acquisition hardware that is not commonly found in hydro plants and is difficult to maintain. Collecting and evaluating data through cavitation surveys to develop CSPs is disruptive to hydro plant operations and data-intensive, especially when developing a cavitation threshold. Many diagnostic methods found in the literature do not suggest a way to establish a cavitation detection threshold, leaving the decision up to the hydro plant operator. A static cavitation detection threshold can become invalid due to changing hydro plant operating conditions, for instance due to: changes in flow rate, hydrostatic head changes (e.g., the reservoir's height changes due to drought or flooding), the number of hydroturbines operating in the hydro plant simultaneously, or disturbances to the inlet or outlet flow. The vibrations that sensors monitor can also be affected by internal changes, causing detection errors from a variety of sources including: repairs made to the hydroturbine runner, worsening of cavitation damage to the runner, faults related to the hydroturbine shaft or bearings, changes in detection instrumentation (intentional or otherwise), and sensor drift. Determining the root cause for a static cavitation detection threshold becoming invalid is difficult, i.e., a stationary threshold cannot determine if plant operating conditions or hydroturbine conditions are the source of the error. Likewise, cavitation intensity measurements are affected by the same problems that affect a static cavitation detection threshold. In summary, existing methods used in industry and in the literature to detect cavitation in a hydroturbine are (1) based on single source measurements, (2) require manual analysis of many different CSPs, or (3) combine many of the same CSPs while experiencing many of the issues noted above.

The first three steps of the RUL prediction process have been carried out in laboratory tests, but the methods used are not practical for monitoring a hydroturbine in a production power plant environment. Complications with data quality, sensor placement, long term robustness of the data collection hardware, and the requirement of manual interaction with the detection system have thwarted attempts to carry out similar tests on production hydroturbines. To our knowledge, results have yet to be published that correlate cavitation erosion rates with data taken from a production hydroturbine. The lack of widespread acceptance or implementation of cavitation monitoring for estimating erosion rates suggests the existing methods are either not effective or not accessible to most hydro-turbine operators.

The issues with establishing a RUL prediction process described above suggest that an adaptive approach that is easily automated would be more successful for long term RUL pre-

diction on production hydroturbines. This paper addresses the first two steps in developing a RUL prediction method: 1) detecting erosive cavitation and, 2) collecting cavitation intensity data. Here cavitation detection is approached by implementing both supervised and unsupervised learning. Cavitation detection is a simple classification problem with two classes: cavitation exists (class 1) or it does not (class -1). With properly labeled training data, many different supervised classification methods can be used to solve this problem. Supervised learning provides a more sophisticated approach to cavitation detection as compared to setting linear thresholds; however, even these algorithms will become inaccurate as sensor data and operating conditions change over time. To solve the problem of drift in the data and operating conditions, a classification algorithm (classifier) could be re-trained at intervals (tantamount to re-calibrating); however, labeled training data would have to be re-generated under the new hydroturbine conditions. The need to manually generate labeled training data impedes the automation of the process and increases the likelihood of miss-classification due to sensor failure, changing operating conditions, or neglect. A more robust approach is to view the creation of training data as an unsupervised learning problem that can be automated once initial parameters are set using domain knowledge. We use this approach to identify operating regions where the hydroturbine is experiencing cavitation using an initially manual process but is then automated to re-calibrate the classifier during ramp-up or ramp-down of the hydroturbine. The intensity of cavitation is determined through calculation of the Mahalanobis distance (MD) from a set of baseline data. The baseline data is generated from the ramp-down or ramp-up data, with the initial ramp-down or ramp-up requiring manual selection of cavitation and cavitation-free operating zones. After the initial manual selection of the operating zones, the process is automated and auto-updates of the cavitation and cavitation-free operating zones are based on the then current hydroturbine running conditions and sensor data.

This paper contributes to the literature a process that addresses the first two steps of developing a RUL prediction for hydroturbines. While the process is demonstrated using proximity probes, it is important to note that this process will work with any sensor commonly used to monitor hydroturbines and is capable of detecting cavitation events. A feature selection method is demonstrated that is simple and can be generalized to many different sensors and CSPs. The feature selection process can be performed on a small amount of data with minimal intrusion to the hydroturbine and hydro plant. After an appropriate CSP is selected, our method can be fully automated, greatly increasing the likelihood of successful long-term cavitation detection and cavitation intensity monitoring. This paper demonstrates using an adaptive threshold that automatically learns the new conditions by collecting a small amount of ramp-up or ramp-down data. We

introduce the MD to hydroturbine cavitation detection and intensity monitoring from the field of cavitation detection in hydraulic pumps where the MD is used as a basis for both establishing cavitation detection thresholds and tracking cavitation intensity. Our method is flexible and multivariate, allowing for the incorporation of many different CSPs thus providing hydro plant operators flexibility in deployment to suit their own specific plant conditions.

1.1. Original Contributions

To summarize, the original contributions of this paper are:

- a method for creating an adaptive cavitation threshold for hydroturbines using machine learning
- a method for automating the detection of cavitation that is appropriate for use in production hydroturbines
- provide necessary tools needed for predicting RUL

2. BACKGROUND

Cavitation is one of the most common faults that occurs in hydroturbines (Dorji & Ghomashchi, 2014; Kumar & Saini, 2010) and the damage caused by cavitation can be very costly to repair (Bourdon, Farhat, Mossoba, & Lavigne, 1999; “The Knowledge Stream - Detecting Cavitation to Protect and Maintain Hydraulic Turbines”, 2014). Cavitation in hydroturbines is the formation of vapor bubbles in the water flowing through the hydroturbine and occurs when abrupt changes in water velocity cause local pressures to fall below the fluid vapor pressure (Dular & Petkovšek, 2015). Vapor bubbles typically develop on or near the hydroturbine runner, but can form in any area where the flowing water reaches higher than expected velocities. When cavitation bubbles collapse, they release large amounts of energy that are destructive to nearby surfaces.

The available water head and flow play a significant role in determining if cavitation will develop during turbine operation (Avellan, 2004). Hydroturbines are designed to prevent cavitation from forming under normal running conditions; however, discussion with hydroturbine operators has revealed several factors outside of the control of designers make eliminating cavitation, and damage caused by cavitation, a difficult task including: 1) available head may change outside of design conditions due to seasonal reservoir variations, floods, or drought; 2) turbulent flow caused by damage or obstructions at the inlet of the hydroturbine; 3) erosion damage on the runner can encourage the formation of cavitation; and 4) the complexity of cavitation formation and collapse makes the amount of damage caused by cavitation difficult to predict in hydroturbines (Dular & Petkovšek, 2015; Jian, Petkovšek, Houlin, Širok, & Dular, 2015).

2.1. Cavitation Detection in Hydroturbines

Hydroturbine researchers generically use the term *cavitation detection* to refer to diagnostic methods that involve sensor measurements, signal processing, and data analysis to aid in determining when cavitation is present (Escaler, Egusquiza, Farhat, Avellan, & Coussirat, 2006; Cencić, Hocevar, & Sirok, 2014; Escaler, Ekanger, Francke, Kjeldsen, & Nielsen, 2014). This definition, however, is ambiguous about key elements of collecting long term cavitation data for studying erosion rates. For the purposes of this paper, we divide cavitation detection into three distinct actions:

- Applying a diagnostic method to sensor measurements to create an indicator sensitive to the onset of cavitation – a CSP as introduced in (McKee et al., 2015).
- Establishing a cavitation threshold (when using a single CSP) or a decision boundary (when using multiple CSPs) that is used to decide when cavitation is present.
- Measuring cavitation intensity in a way that can be used to calculate or estimate cavitation erosion rates.

Many diagnostic methods are available to hydroturbine operators for creating a CSP (Varga, JJ and Sebestyen, Gy and Fay, 1969; Bajic, 2002; Escaler et al., 2006; Rus, Dular, Sirok, Hocevar, & Kern, 2007; Cencić et al., 2014; Escaler et al., 2014). Unfortunately, cavitation intensity measurements are not directly addressed in these diagnostic methods and the action of establishing a cavitation threshold is completely ignored. With the exception of work by Dorey, et al. (Dorey, J.M.; Laperrousaz, E.; Avellan, F.; Dupont, P.; Simoneau, R.; Bourdon, 1996), in collaboration with Bourdon, et al. (Bourdon et al., 1996) and continued by Francois (Francois, 2012), cavitation diagnostics studies have focused on short term data collection and manual data analysis.

2.2. Instrumentation for Cavitation Detection

When a cavitation bubble collapses on the surface of the hydroturbine runner, the shock wave it creates propagates through the hydroturbine and surrounding water. Cavitation creates significant erosive damage when thousands of bubbles collapse over a short period of time producing vibration response between 3000 and 400,000 Hz (Escaler et al., 2006; Cencić et al., 2014). Detecting the high frequency response of cavitation directly requires sophisticated sensors and equipment meant for high frequency applications, thus accelerometers and acoustic emission sensors are frequently used.

Since hydroturbines have relatively low shaft speeds (typically well below 20 Hz (Gordon, 2001; Escaler et al., 2006), high frequency monitoring equipment is specific to cavitation detection. Other fault conditions such as balance and alignment problems occur at frequencies below 500 Hz and are monitored with low sample rate data acquisition equipment

and proximity probes that produce a signal proportional to the relative movement between the sensor and the hydroturbine shaft. Due to added cost, more sophisticated cavitation detection sensing is not typically included on production hydroturbines.

Pennacchi, et al. (Pennacchi, Borghesani, & Chatterton, 2015) showed that proximity probes can be used for diagnosing cavitation. Instead of measuring cavitation events directly, they used synchronous averaging and spectral kurtosis to monitor the hydroturbine shaft's natural frequency response fluid instability. In their implementation, the signal is filtered around the natural frequency of the shaft.

2.3. Cavitation Intensity

Dular et al. (Dular, Stoffel, & Širok, 2006) developed a cavitation damage model that expresses cavitation damage, A_{rel} (damage area), in terms of the time a surface is exposed to cavitation, τ , the cavitation shedding frequency, f , the probability of a cavitation event (referred to as a micro-jet by Dular et al.), $P(mj)$, and the velocity characteristics of the flowing water, v_{ref} and v , as follows:

$$A_{rel} = \frac{A_{pit}\tau f P(mj)}{A_{ref}} \left(\frac{v}{v_{ref}} \right)^2 \quad (1)$$

where A_{ref} is the total reference area and A_{pit} is the pit area. The damage model was verified on a radial pump with f and v being measured during the experiment and $P(mj)$ being held constant. The significance of this model is that cavitation damage is related to cavitation intensity based on local fluid velocity, exposure time, and the frequency of cavitation events.

In a practical implementation one must choose sensor types and locations as well as CSPs that give reliable intensity measurements. Variation in the structure and layout of different hydroturbines combined with different sensor types and placement make amplitude measurements difficult to compare. The measurement scale (or unit) of a CSP is dependent on the sensor type and the measured value is affected by the sensor location (Schmidt et al., 2014). Cavitation tests on production hydroturbines are usually performed with accelerometers and acoustic emission sensors placed on the upper and lower hydroturbine bearings as well as the stems of the guide vanes that control water flow rate into the turbine runner (Bajic, Services, Gmbh, & Zithe, 2003; Escaler et al., 2006; Cencic et al., 2014; Escaler et al., 2014). Proximity probes are typically located in or near the hydroturbine's bearings. Each accelerometer, acoustic emission sensor, and proximity probe will produce a signal with a different amplitude. Unfortunately, this means cavitation intensity measurements gathered directly from the sensor's native measurement scale can only be performed once the sensor's response to cavitation excitation is known.

To address the issue of signal amplitude variation, data normalization is used. Z-score standardization is a popular method of normalization when comparing and analyzing multivariate data with different amplitude scales (Milligan & Cooper, 1988; Keogh & Kasetty, 2002; Nandi, Liu, & Wong, n.d.; Kan, Tan, & Mathew, 2015). Z-score standardization - often called 'standardization' - linearly transforms the data to have a mean of zero and a variance of 1. A data set $X = [x_1, x_2, \dots, x_n]$ is standardized by normalizing the difference between the set mean μ_x and each set value by the set standard deviation, σ_x , as shown below:

$$\hat{X} = \frac{x_i - \mu_x}{\sigma_x} \quad \text{for } i = 1 \dots n \quad (2)$$

The standardized amplitude values are unit-less and measure the distance, in standard deviations, from the mean of the data. In vibration analysis, standardization prevents high amplitude signals from dominating the analysis and obscuring important low amplitude features.

Standardization is frequently used as a data preparation step for machinery diagnostics and prognostics (Saxena, Celaya, Saha, Saha, & Goebel, 2009; Khelf, Laouar, Bouchelaghem, Rémond, & Saad, 2013; Ramasso & Saxena, 2014; Kan et al., 2015); however, we were unable to find it as a step in any published hydroturbine cavitation diagnostic research. Instead of standardization, researchers apply other methods of normalization such as dividing a set of frequency spectra by the first spectrum collected (Bajic, 2002; Cencic et al., 2014) or do not normalize at all. Presumably, normalization is not deemed necessary because researchers and practitioners often compare vibration signals that have the same magnitude scale or are following a collection and analysis process specified in an international standard (ISO, 2005). We choose to standardize our vibration signals for two reasons: 1) vibration amplitude has a non-linear relationship with respect to frequency (e.g., acceleration scales with the square of the frequency relative to displacement, $a = 2df^2$), and 2) vibration amplitude is affected by the transmissibility between the vibration source and the sensor location, i.e., sensors installed at different locations will observe different amplitudes for the same vibration event (Schmidt et al., 2014). We have found that standardizing signals between different types of sensors, sensor locations, and frequency ranges allows for a consistent comparison of vibration amplitude.

2.4. Mahalanobis Distance

Cavitation detection can be viewed as an online process that examines new vibration signal observations, i.e., x_{n+1} , to determine if cavitation is present. By using a baseline of vibration data when no cavitation occurs to determine a μ_{base} and applying the concept of standardization as expressed in Eq. 2 one can assess the difference between the current reading and

the baseline measurement. The Mahalanobis distance, Eq. (3) is a multivariate extension of this concept that is useful for outlier detection, structural health monitoring, clustering, and detecting cavitation in pumps (De Maesschalck, Jouan-Rimbaud, & Massart, 2000; Figueiredo, Park, Farinholt, Farrar, & Lee, 2012; Inacio, Lemos, & Caminhas, 2014; McKee et al., 2015).

$$MD = [(x_{n+1} - \mu_{base})\Sigma^{-1}(x_{n+1} - \mu_{base})^T]^{\frac{1}{2}} \quad (3)$$

In the multivariate case, X now becomes a set of variables, such as observations from multiple sensors while the hydro-turbine is in a healthy state, and x_{n+1} contains the next observation from every sensor. The covariance matrix, Σ , is calculated using the expression:

$$\Sigma_x = \frac{1}{n-1}(X)^T(X) \quad (4)$$

The Mahalanobis distance (MD) is useful for cavitation detection because it takes into account the correlation of the sensor data and allows us to describe and compare the distribution of several sensors using a single metric. In terms of establishing a threshold for identifying cavitation, instead of creating a threshold for each available sensor, we now have a single threshold that incorporates all the signals. When X contains observations from a single sensor, the MD reduces to an expression similar to Eq. (2), i.e.,

$$MD = \left[\frac{(x_{n+1} - \mu_{base})^2}{\sigma_{base}^2} \right]^{\frac{1}{2}} \quad (5)$$

This single variable form no longer contains a covariance matrix, but still takes into account the distribution of the healthy baseline data for its distance metric. Equation (5) should be used when only one sensor is available for cavitation measurements, or when sensor signals are modeled as independent observations.

2.5. Prognostics and Erosion Rate Prediction

Prognostics can be defined as the process of forecasting the remaining useful life RUL, probability of failure, or future condition of a component or system (Jardine, Lin, & Banjevic, 2006; An, Kim, & Choi, 2013; Kan et al., 2015). Prognostic models are categorized as physics-based, data-driven, or combination approaches. Physics-based models require a mathematical understanding of the degradation phenomenon affecting the system of interest, whereas data-driven models rely on condition monitoring or training data collected from the system. Under the right circumstance, both models are effective. In practice, both strategies are needed since mathematical models require experimental validation, which is fundamentally data driven. Similarly, data-driven meth-

ods require an understanding of the underlying physics to collect meaningful data. Current physics-based approaches for cavitation prognostics focus on predicting *erosion rates*. The underlying mechanisms of cavitation have been shown to be quite complex (Dular & Petkovšek, 2015), yet numerical methods developed for erosion rate prediction have been experimentally verified in simplified systems (Flageul, a Archer, & C, 2012; Jian et al., 2015). Though progressing, numerical methods for predicting erosion rates have yet to be verified under conditions and geometries as complex as an operating hydroturbine. Physics-based prognostic models require knowledge of very complex environments and mechanisms that make them hard to build for practical applications (Heng, Zhang, Tan, & Mathew, 2009; Kan et al., 2015).

Researchers developing data-driven prognostic models also focus on estimating erosion rates. As previously mentioned, laboratory experiments have verified that damage caused by cavitation is related to cavitation intensity, which in turn can be measured through vibration and acoustic emission. Producing similar results outside of the controlled environment of the laboratory has proven to be much more complex. Hammitt and De discussed predicting erosion rates from sensor measurements as early as 1979 (Hammitt & De, 1979), but focused primarily on cavitation erosion on simple shapes in laboratory environments. Francois (Francois, 2012) has written about a major power producers' attempts at erosion rate estimation; however, no results have been published as of yet. Wolff, Jones and March (Wolff, Jones, & March, 2005) attempted a similar endeavor at another major power plant in an attempt to establish an erosion rate model, but insufficient data stymied this effort. Similar research in other fields has shown that data-driven prognostic models are often plagued by problems with data quality and data quantity. It is for this reason that we focus our research in this paper on improving long-term cavitation detection and intensity monitoring for production hydro plants.

3. METHODOLOGY

In this section, we present a methodology for collecting the sensor data needed to create remaining useful life models for hydroturbine runners. The underlying concept of our methodology is that sensor signals collected from a hydroturbine ramp-down and ramp-up (a small data set that requires minimal disruption to power production) can be used to 1) select a CSP, 2) create a threshold for identifying cavitation, and 3) create a baseline for measuring cavitation intensity. When automated means are used for creating training sets (an unsupervised learning problem) and for classifying cavitation (a supervised learning problem), our method can be used to create a fully automated cavitation detection strategy that can adjust for sensor drift and changes in operating conditions of the hydroturbine.

We approach cavitation detection from a machine learning framework by breaking it into four steps: 1) Select Cavitation Features, 2) Create Training Sets, 3) Train Classifier, and 4) Measure Intensity. Our methodology was developed using vibration data collected from four proximity probes mounted on an 85 MW hydroturbine. Our feature selection process can easily be used with other sensor types more commonly selected for cavitation detection including accelerometers or acoustic emission sensors; however, an advantage to using proximity probes for cavitation detection is that many older hydroturbine units are permanently instrumented with proximity probes. This is often not the case for accelerometers and acoustic emission sensors that have higher frequency response ranges, but require hardware capable of faster sampling rates. Additionally, the use of four sensors demonstrates how the method has multi-dimensional capability that both improves the classification accuracy and is more robust for longer term usage since it doesn't rely on a single signal source that can more easily be corrupted by noise.

3.1. Select Cavitation Features

In this work, the feature being selected is the frequency range used for the CSP calculations used to predict when a hydroturbine is experiencing cavitation. This definition for a feature could easily be expanded to include the sensor type and sensor location when these additional options exist (Gregg, Steele, & Bossuyt, 2016).

Step 1: Collect Ramp-Down Data

The features used in our method are created from raw data collected from the hydroturbine as it ramps linearly between its maximum and minimum power output running conditions¹. When using proximity probes for cavitation detection, the minimum sampling rate used to collect the data should be based on the higher of either the blade passing frequency, f_b , or the guide vane passing frequency, f_v . For a give hydroturbine running speed, N , f_b and f_v are defined as follows: $f_b = N(\# \text{ of runner blades})$, and $f_v = N(\# \text{ of guide vanes})$.

Based on the typical values of running speed, the number of guide vanes, and the number of runner blades on hydroturbines found in literature (Escaler et al., 2006; Cencic et al., 2014), and taking into account the Nyquist theorem, a sample rate of at least 1,000 Hz is recommended.

The amount of time in seconds the hydroturbine takes to go through the ramp-down will affect the amount of data collected, its frequency resolution, and total number of points available to create training data. Here a 60 - 90 second ramp-down produces sufficient data, however, these lengths were

based on the data available for our analysis.

Step 2: Calculate the Variance of Each Frequency

In Step 2, we search for vibration frequency ranges in the ramp-down data that significantly change in amplitude over time. During the hydroturbine ramp-down, the speed of the turbine remains constant and the only variables that change are *generation load and water flow* through the turbine. Vibration frequencies dependent on water flow can be further analyzed to determine if they are related to cavitation. The following process, when applied to the ramp-down data collected in Step 1, allows us to identify frequencies dependent on water flow: 1) the ramp-down data is divided into 1 second blocks, 2) the direct current (DC) (zero frequency) trend is removed in each block resulting in data centered around zero, 3) the discrete Fourier transform (DFT) of each block is computed, and 4) the sample variance of each frequency value across all blocks is calculated.

The frequency resolution of a spectrum, f_{res} , is dependent on the period of the data collected, T , and correspondingly, the sample frequency, f_s , and the number of data samples, N .

$$f_{res} = \frac{1}{T} = \frac{f_s}{N} \quad (6)$$

By selecting a ramp-down data block length of 1 second, the resulting DFT calculation will produce a spectrum with a resolution of 1 Hz, which is sufficient to differentiate cavitation related frequencies ranges within the ramp-down data. The total number of 1 second blocks of data that will be created, t , is dependent on the total length of ramp-down data collected. Selecting block lengths of 1 second provides both sufficient frequency resolution and training data.

When used to detect shaft vibration on a hydroturbine, proximity probes produce a signal proportional to the distance between the tip of the proximity probe and the surface of the turbine shaft. The vibration signal from a proximity probe will therefore oscillate around the average distance between the proximity probe and the shaft which adds a DC offset to the signal. In addition to the added offset, each vibration block is likely to have a slight linear trend in the DC portion of the signal that will cause the DFT to have a large zero frequency amplitude that obscures the amplitude of higher frequencies of interest. The DC offset and linear trend should be calculated and subtracted from each data block. The DFT of each block of ramp-down data can then be calculated using the fast Fourier transform algorithm (Cooley & Tukey, 1964).

Recall that the flow rate of water through the turbine runner is the only running condition variable that changes in the hydroturbine during ramp-down. As noted by Escaler et al. (Escaler & Egusquiza, 2003; Escaler et al., 2006, 2014), cavitation is related to flow rate and causes vibration at multiple

¹The direction of the ramp – up from minimum to maximum or down from maximum to minimum power output – is not important to the research presented here, although in other applications the differences in ramp – up and ramp – down are important. Throughout this text, unless otherwise noted, we generically use the term ramp-down to signify a ramp in either direction.

frequencies including running speed, f_b , and f_v , as well as through broad-band high frequency noise. As such, vibration frequencies with significant change in amplitude throughout the ramp-down data are marked as being related to cavitation. The change in amplitude of vibration frequencies throughout the ramp-down is expressed by the variance of each column of the \hat{Y} matrix. Variance is calculated from the mean, μ , using (Montgomery & Runger, 2007).

$$var = \frac{\sum_0^n [x_i - \mu]^2}{n - 1} \quad (7)$$

The result of applying Eq. (7) to the columns of \hat{Y} is a single vector that is plotted to form a variance frequency spectrum. The variance frequency spectrum is used to quickly identify frequencies that change during ramp-down and are subsequently related to changes in water flow rate through the hydroturbine.

Step 3: Select CSP Frequency Ranges

The CSP chosen in our methodology for cavitation detection is calculated from the root mean square (RMS) amplitude of proximity probe vibration within one or multiple frequency bands. CSPs based on RMS calculations and frequency filters have been shown to be effective for cavitation detection in both hydroturbines and pumps, (McKee et al., 2015; Cencic et al., 2014) and are practical to implement since they can be easily derived using either digital or analog methods²

The frequency bands to use for RMS calculations are based on the variance frequency spectrum created in Step 2. More generally, when using proximity probes for cavitation detection in hydroturbines, three frequency regions are of interest:

1. Vibration frequencies below running speed are affected by draft tube swirl, and Von Karmen vortex shedding, or other hydraulic instabilities (Escaler et al., 2006).
2. Increased vibration frequencies at running speed can also be an indicator of hydraulic instability; however, running speed vibration may also be influenced by other types of faults including unbalance, misalignment, and bearing wear. (Egusquiza, Valero, Valentin, Presas, & Rodriguez, 2015).
3. High frequency vibration at f_v , f_b , as well as general broadband vibration is associated with cavitation that causes erosion on runner blades.

3.2. Create Training Sets

In our methodology, we treat erosive cavitation detection as a binary classification problem with categories: CAVITATION and NO-CAVITATION, and numerically represent them as

$\{1, -1\}$, respectively. For reasons described earlier, we use MD to establish labels for the initial set of training data. Standardizing MD helps with separation of data and interpreting the results. Each point in the training set can be categorized manually, or in an automated fashion using an unsupervised learning algorithm using these steps:

1. Band pass filter the previously collected sensor data ramp-down signals around each frequency range of interest determined from the Cavitation Feature Selection step.
2. Divide the filtered signals into 1 second blocks and calculate the RMS of each block. The result is a ramp-down data set for each frequency range of interest ($x_{f1} \dots x_{fn}$).
3. Select the baseline data for calculating MD by plotting the standardized RMS amplitude of each ramp-down data set versus sample number and identifying a continuous sample range free from cavitation or other faults. This sample range, $X_{baseline}$, is the baseline data and is meant to be representative of the fault free distribution of the data for each frequency range and sensor. As a general rule of thumb, the baseline data should contain at least 30 samples (Montgomery & Runger, 2007)³.
4. Combine the ramp down sets, ($x_{f1} \dots x_{fn}$), into a single matrix X and calculate the MD of the values in the ramp-down data sets by first calculating the covariance matrix of $X_{baseline}$ with Eq. (4) then applying Eq. (3) to the remaining values in X . Values for μ are calculated from $X_{baseline}$. The MD values can then be standardized by applying Eq. (2).
5. To categorize the data manually, use ramp-down data from the frequency range(s) most representative of erosive cavitation and plot the standardized MD of the ramp-down data versus sample number. Select a cavitation threshold value that is visually above the points in the $X_{baseline}$ sample range. When using standardized MD, a conservative threshold, corresponding to fewer false positives, will be close to 1 and a more aggressive threshold, corresponding to more false negatives, will be close to or below 0. All points with a MD larger than the threshold belong in the CAVITATION category and all the other points belong in the NO-CAVITATION category.
6. To automate data categorization, instead of visually selecting a threshold, use an unsupervised learning algorithm such as k-means clustering (Pollard, 1981) to separate the ramp-down data from Step 5 into two clusters. The NO-CAVITATION cluster should minimally contain all the samples in the $X_{baseline}$ range.

²For accelerometers or acoustic emission sensors, demodulation methods (Bourdon, Simoneau, & Lavigne, 1989; Escaler et al., 2006, 2014) can also be used as a basis for the CSP; however, demodulation relies on the use of frequency bands beyond the sensing capabilities of proximity probes.

³If the RMS values of the vibration data is assumed to be normally distributed, the number of sample points can be reduced (Toyota, Niho, & Chen, 2000).

3.3. Train a Classifier

Once cavitation features are selected and training sets are created, cavitation detection is automated by applying a classification algorithm to new cavitation features that are generated to predict if the hydroturbine is experiencing erosive cavitation. When classifying cavitation with a supervised machine learning algorithm, an additional training step is required that allows the algorithm to generate its own cavitation threshold (more generally, this is called a decision boundary or hyper-plane) from the training sets created in the previous step.

As a method for evaluating classifiers, we suggest comparing the classifier predictions to a naïve, single variable algorithm that calculates the standardized MD, \hat{x}_{MD} , of each new value based $X_{baseline}$ and compares this value to the threshold established to create the training set. Given a threshold, pseudocode for this classifier is as follows:

```

FOR new RMS value  $x$ 
  calculate  $\hat{x}_{MD}$ 
  IF  $\hat{x}_{MD} > threshold$ 
    classify  $x$  as 1 (Cavitation)
    Calculate and save cavitation intensity
    based on  $x_{MD}$ 
  ELSE
    classify  $x$  as -1 (No Cavitation)
END

```

The accuracy obtained by applying the naïve classification algorithm can be used as a baseline for comparing more sophisticated classification algorithms. The advantages of using a naïve classifier are ease of implementation, low computing cost, which makes it feasible to use in either an on-line or batch mode, and good accuracy. The disadvantage of such a simple classifier is that it is based on a single variable that is not sensitive to other, non-cavitation related faults so it cannot be used for more generalized fault detection. A multi-dimensional classification algorithm such as a support vector machine (SVM) may be used to take advantage of features created from other frequency ranges to both enhance cavitation detection and classify other fault states such as non-erosive cavitation.

3.4. Measuring Cavitation Intensity

We use the MD of the CSP most representative of erosive cavitation as our cavitation intensity measurement. MD is suited well for measuring cavitation intensity because it automatically accounts for variability in the sensor signal. The benefit of this is best shown graphically using real hydroturbine data. Figure 1 shows RMS vibration amplitude with respect to time of a hydroturbine going through a ramp-down as measured by two sensors mounted at different locations. Sensor 1 clearly

records a higher maximum amplitude as well as accumulated amplitude (area under the curve) from the erosive cavitation zone in sample range 11 - 38. It is also evident that Sensor 1 increases in amplitude more than Sensor 2 over the baseline range from sample 55 - 100. By contrast, Figure 2 shows the same sensor data, but with amplitude measured as MD. Sensor 2 now clearly shows a higher total, and accumulated amplitude since the MD calculation takes into account the lower variance (as measured by standard deviation) of Sensor 2 through the base-line range. In this way, signals that are more stable when cavitation is not present can contribute more to the intensity measurement.

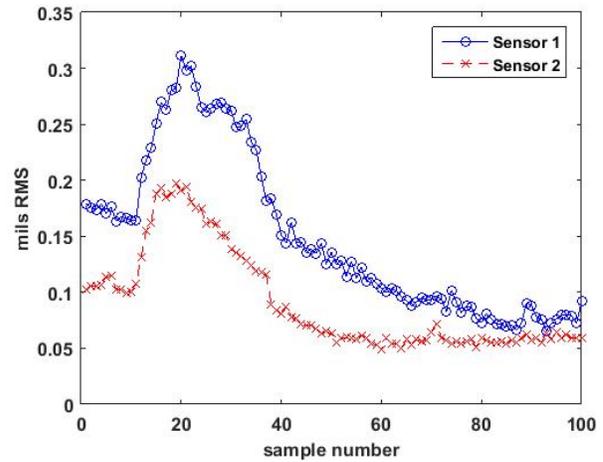


Figure 1. Sensor vibration amplitude comparison from a hydroturbine ramp-down as measured in RMS

In summary, the first half of our cavitation detection method uses hydroturbine ramp-up or ramp-down data as an input and returns a training set of CSPs that are used to establish erosive cavitation thresholds. The second half of our cavitation detection strategy creates cavitation thresholds based on the CSPs as described above. The final output is a set of cavitation CSPs created from the sensors being used for cavitation detection, and cavitation thresholds that can adapt to changes in running condition whenever the hydroturbine goes through a ramp-up or ramp-down.

4. CASE STUDY

We present here a case study using vibration data collected from an 85 Megawatts (MW) hydroturbine known to be experiencing erosive cavitation and located at a hydro power plant in the American West⁴. Vibration data were collected from four proximity probes mounted 90 degrees apart facing the hydroturbine's main shaft. Proximity Probes 1 and 2 were located near the hydroturbine's lower bearing while Proximity

⁴Our data source has asked for the exact location and details of the hydro plant to remain confidential.

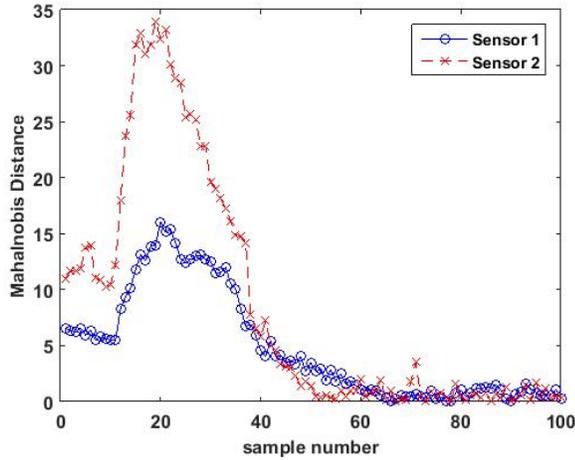


Figure 2. Sensor vibration amplitude comparison from a hydroturbine ramp-down as measured in Mahalanobis Distance

Probes 3 and 4 were located near the upper bearing. Signals from the proximity probes were sampled at a rate of 10,000 Hz. The data used for feature selection and to create training sets were collected while the hydroturbine ran through a continuous ramp-down from 85 MW to 0 MW over a 100 second period, which was divided into 1 second blocks. The power produced versus time by the hydroturbine during the ramp-down is shown in Figure 3.

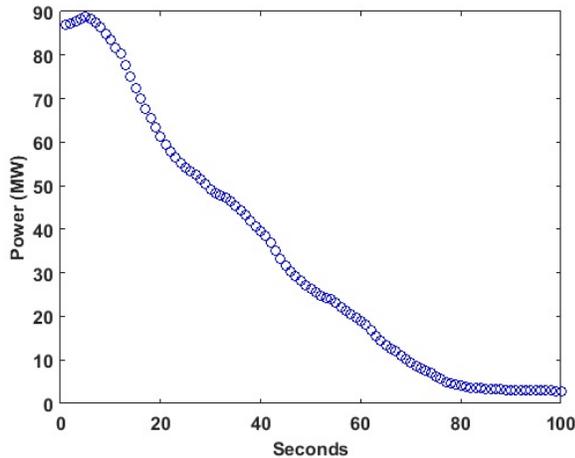


Figure 3. Hydroturbine power versus time during the ramp-down

The goal of this case study is to both demonstrate the methodology presented in this paper and compare hydroturbine cavitation classification accuracy using the following four approaches: 1) Classify cavitation with a naïve threshold classifier and a manually selected cavitation threshold, 2) Classify cavitation with a naïve threshold classifier and a cavitation threshold found by applying an unsupervised learning algo-

rithm, 3) Classify cavitation with a supervised learning algorithm and training data that is manually labeled, and 4) Classify cavitation with a supervised learning algorithm and training data that is labeled by applying an unsupervised learning algorithm. A SVM was selected as the supervised learning algorithm to use for predicting cavitation classes and a K-Means was selected as the unsupervised algorithm for labeling training data. SVM and K-means algorithms used for this case study are based on the corresponding built-in functions of Matlab (v2015a) with the Statistics and Machine Learning Toolbox.

The SVM, as described by Cortes and Vapnik (Cortes & Vapnik, 1995), is a machine learning algorithm for binary classification problems that is frequently used to detect machine faults in the field of condition monitoring (Widodo & Yang, 2007). SVMs were selected for this case study due to their high accuracy, low computational burden, ease of use, and popularity in the machine learning community (Samanta, Al-Balushi, & Al-Araimi, 2003; Witten & Frank, 2005; Wu et al., 2008).

K-Means clustering, described by Hartigan (Hartigan, 1975), is a heuristic algorithm that aims to divide M data points into K clusters so that the sum of squares is minimized within each cluster. The K-means algorithm used in this case study (Lloyd, 1982; Hartigan & Wong, 1979) is iterative and requires the practitioner to choose a value for K as well as K data points, called seeds, that are initially assigned to their own cluster. Next, the point to cluster centroid distance of each data point is calculated and all points included in the cluster analysis are assigned to the cluster with the closest centroid. The new cluster centroid is then calculated and the data points are then re-assigned based on the new centroid. This repeats until clusters are no longer re-assigned after the new centroids are calculated. The final cluster results are dependent on the value of the K seeds selected for the first centroid calculation. To obtain consistent results for establishing a cavitation threshold, a segmentation technique similar to bi-level thresholding (Pal & Pal, 1993) was used where the input value of K was always equal to 2, and the minimum and maximum CSP values in the training set were used as seeds.

4.1. Step 1: Select Cavitation Features

A fast Fourier transform (FFT) was calculated for each block of ramp-down data, then the variance spectrum was created to determine how vibration frequencies responded during the hydroturbine ramp-down. The proximity probes responses to the ramp-down as can be seen in Figure 4. Based on the variance spectra, three frequency ranges (Figure 5) were identified as features to use for calculating CSPs:

$$\text{Frequency Range 1} = 1 - 3 \text{ Hz}$$

$$\text{Frequency Range 2} = 3 - 30 \text{ Hz}$$

$$\text{Frequency Range 3} = 50 - 90 \text{ Hz}$$

Frequency Range 1 is made up of frequencies below running speed, while Frequency Range 2 includes the shaft rotating frequency and its first several harmonics. Frequency Range 3 includes the hydroturbine blade-pass and vane-pass frequencies.

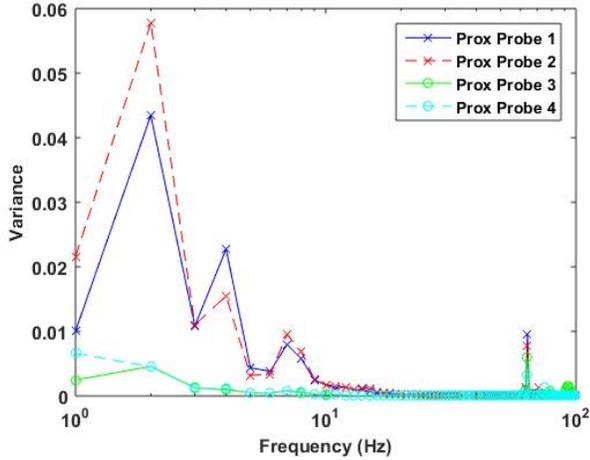


Figure 4. Variance spectrum of all four proximity probe signals

4.2. Step 2: Create Training Sets

Next, the sample range to use as baseline data for the MD calculation, $X_{baseline}$, was chosen by analyzing the standardized amplitude of the three CSPs in the time domain over the hydroturbine ramp-down as shown in Figure 6. Because the CSP's amplitudes for samples 55 to 100 are relatively low and steady, this sample range was chosen for $X_{baseline}$.

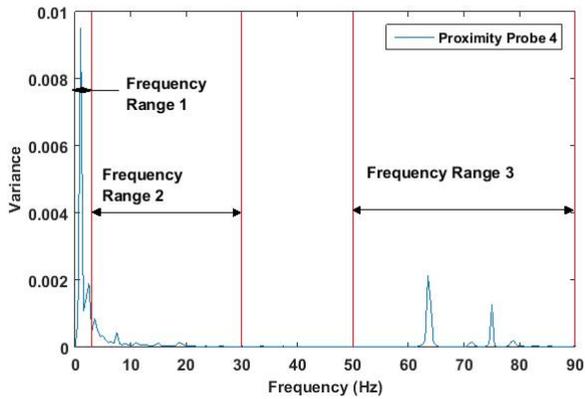


Figure 5. Variance spectrum showing the frequency ranges used for calculating the CSP values

As previously explained, Frequency Range 3 is expected to be the most sensitive to erosive cavitation; however, when a multi-dimensional classifier is used for prediction, all three ranges can be used to improve accuracy. One reason for the

improvement in accuracy is that each frequency range has an independent response to flow during the ramp-down. The independence of each CSP is evident when comparing their standardized amplitude in the time domain during the ramp-down, as shown in Figure 6.

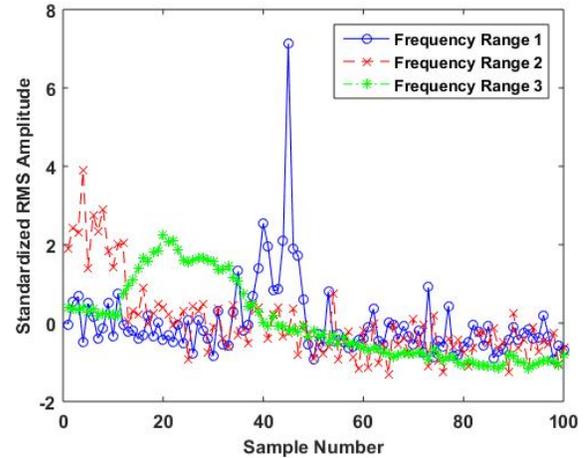


Figure 6. CSP values plotted versus ramp-down sample number

Once $X_{baseline}$ was selected, the standardized MD distance was calculated for all of X . It is important to note that MD can be calculated in its multivariate form, Eq. (3), where X is a combination of CSPs from all the proximity probes, or the single variable MD can be calculated independently for each sensor using Eq. (5). When performing the multivariate calculation, there will be a single set of MD values which means only a single threshold will be needed for all the sensor measurements. However, the single variable calculation will produce 4 sets of MD values and 4 thresholds, only one of which will need to be selected for labeling training data. Results from both methods are presented in our case study.

Based on the Frequency Range 3 CSP values, cavitation thresholds were first selected manually – using both the multivariate MD calculation and the single variable method – then by automating the method utilizing a K-means clustering algorithm. The selected cavitation thresholds are shown in Table 1.

These cavitation thresholds are used for labeling training sets as well as for classifying cavitation when applying the naïve classifier. Training sets for binary classification can only have one label; however, a unique set of labels will be produced for each proximity probe due to slight variations in amplitude between each sensor. For example, Figure 7 shows several CSP values between sample 1 and sample 10 are above the cavitation threshold for Proximity Probe 3, but below the cavitation threshold for the other proximity probes. In our analysis, the classification labels established by applying the thresholds to

Single Variable Thresholds	
Manual 1	0.018
Manual 2	-0.110
Manual 3	-0.386
Manual 4	-0.392
K-means 1	0.338
K-means 2	0.162
K-means 3	0.494
K-means 4	0.622
Multivariate Thresholds	
Manual	0.474
K-means	0.380

Table 1. Cavitation thresholds for labeling training data

data from Proximity Probe 2 were used for labeling the training sets. Figure 8 shows the multivariate threshold and resulting classification labels found by applying a K-Means clustering algorithm to the training set containing all the proximity probe CSPs.

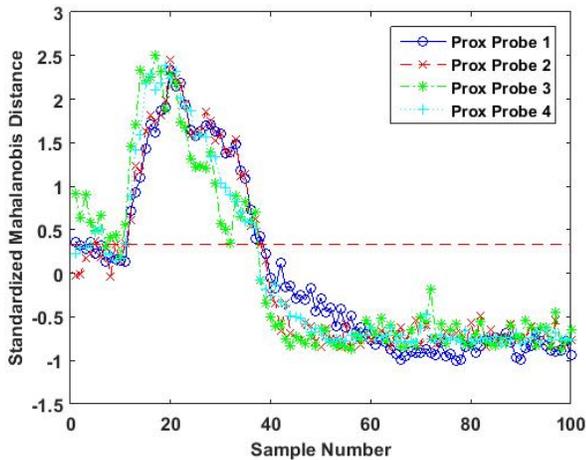


Figure 7. The manually selected, single variable cavitation threshold (dashed red line)

4.3. Step 3: Train Classifiers

The naïve classifier does not require additional training beyond establishing a cavitation threshold. A label for each new observation is generated by directly comparing its standardized MD to the cavitation threshold then labeling the observation "1" if the value is above the threshold, or "-1" if it is not. The observations the naïve classifier uses for comparison are data points from one proximity probe, for the single variable case, or all the proximity probes, for the multi-

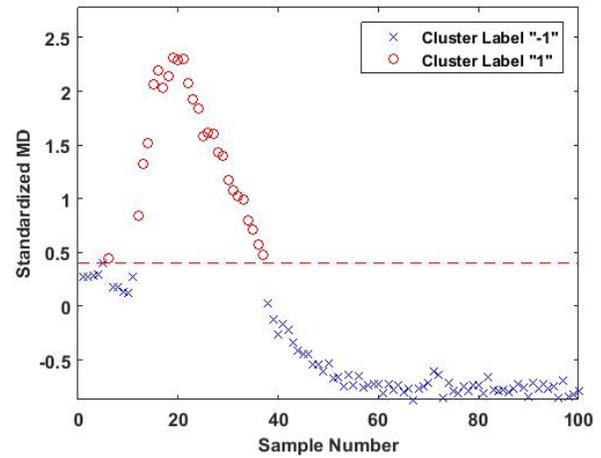


Figure 8. The multivariate cavitation threshold found through k-means clustering (dashed red line)

variate case, with a single value based on the CSP calculated from Frequency Range 3. In other words, the naïve classifier uses a one-dimensional cavitation threshold and acts on one-dimensional data. Accuracy testing for the naïve classifier included one test for each proximity probe, and one for the multivariate threshold, using both the manually and K-Means selected thresholds, resulting in 10 accuracy values.

The SVMs also rely on the labeled training sets to construct a decision boundary; however, the boundary can be multi-dimensional, which means the training set and testing set can simultaneously include any or all of the proximity probes and CSPs. The benefits of a multi-dimensional decision boundary include more accurate classification predictions on data that is not linearly separable as well as the ability to extend the capabilities of a classifier to recognize more than just two categories of data. The multi-dimensional capability of a SVM also means a decision must be made about which proximity probes and CSPs to include in the training. For our analysis, we decided to train and test a SVM for every unique combination proximity probe and CSP, and compare the combinations with the highest accuracy. There are 4 proximity probes, and 3 CSPs for each of the proximity probes, which means that there are 12 individual training sets and $2^{12} - 1 = 4095$ unique combinations of these 12 training sets.

We also looked at the multivariate threshold case where there is only one CSP for each frequency range for a total of 7 unique combinations. A potential advantage of using an SVM is its capability to find non-linear thresholds. The correctly classified test data (Figure 7) shows that a non-linear cavitation threshold may be appropriate. To test this hypothesis we trained SVM models with polynomial kernels with orders 1 to 8 to test how a non-linear boundary affected classification accuracy. Non-linear SVM models were only trained and tested for the multivariate threshold case.

4.4. Step 4: Verification - Classification Test Results

Data used for testing accuracy of the SVM and naïve classifiers as well as calculating cavitation intensity were collected while the hydroturbine ran for prolonged periods in 17 unique flow rates ranging from 5 MW to 85 MW in 5 MW increments; 24 seconds of data was collected for each flow rate which was then divided into 1 second blocks resulting in 408 total blocks of vibration data used to create the test data. Other running condition variables such as hydrostatic head, other turbines in the plant operating, and other factors were held effectively constant throughout the data collection period. The correct class labels for the training set were created manually using more traditional cavitation detection methods as well as sensor data from accelerometers and acoustic emission sensors. Additional information on the full analysis and general cavitation detection methods used to create the class labels can be found in (Gregg et al., 2016; US Department of the Interior Bureau of Reclamation, 2014; Escaler et al., 2006; Escaler & Egusquiza, 2003).

The naïve and SVM classifier algorithms were applied to the test data and the resulting class predictions were compared to the correct class labels to determine the prediction accuracy. Cavitation intensity was calculated directly from the MD of the test data. The accumulated cavitation intensity over the whole data set, I_{total} , is calculated by taking the MD of each CSP identified by the classifier as being in the cavitation class, $X_{MD-cavitation}$ and multiplying it by the time block length used to create the CSP, t_{block} as shown in Eq. (8). For the training and testing data, the time block length is 1 second and only CSPs created from Frequency Range 3 are used for intensity measurements.

$$I_{total} = \sum (X_{MD-cavitation})(t_{block}) \quad (8)$$

Classifier accuracy results for the top performing training set combinations based on single variable thresholds are shown in Table 2 (see Appendix). For the SVM results, proximity probe/CSP pairs are abbreviated with the proximity probe number first, "-", then 'CSP' followed by the frequency range used to create the CSP. For example, a training set created with data collected from Proximity Probe 1 that uses Frequency Range 1 for the CSP calculation would be abbreviated "PP1-CSP1".

Classifier accuracy results based on multivariate thresholds are shown in Table 3(see Appendix). Since all the proximity probe data are combined in the multivariate case, only the frequency range used for training and the order of the non-linear polynomial threshold are noted.

Figure 9 graphically shows the correct classification labels for the test data. Labels predicted by the naïve classifiers are shown in Figures 10 and 11. Labels predicted by the SVM

classifiers are shown in Figures 12 and 13.

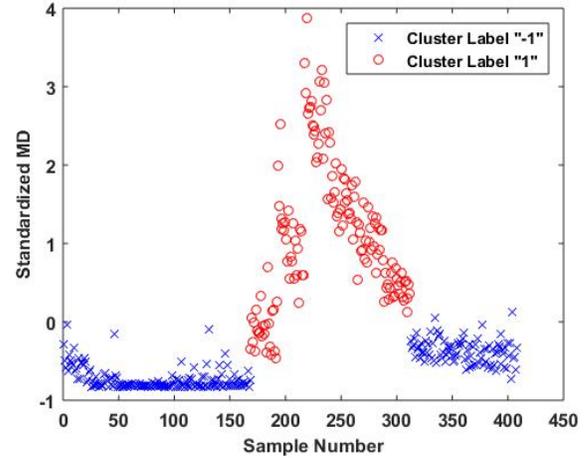


Figure 9. Test data shown with correct classifications using traditional, manual analysis techniques

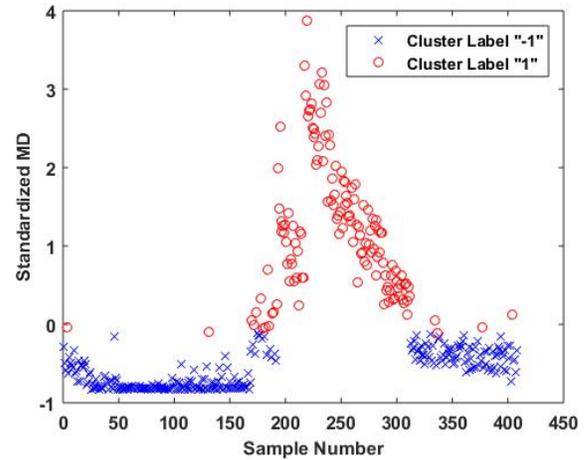


Figure 10. Test data shown with labels predicted by the naïve classifier using manually selected threshold

5. DISCUSSION

The methodology outlined in this paper provides several benefits when compared to other cavitation detection strategies. Additionally, our method addresses common problems associated with cavitation thresholds and intensity measurements. While the method presented here does not yet provide cavitation erosion rate calculations, it provides the tools necessary to automate the collect of cavitation intensity data, a crucial step toward creating an erosion rate model for production hydroturbines.

The cavitation detection process described in this paper was demonstrated on ramp-down data collected from proximity probes on a hydroturbine experiencing erosive cavitation. This

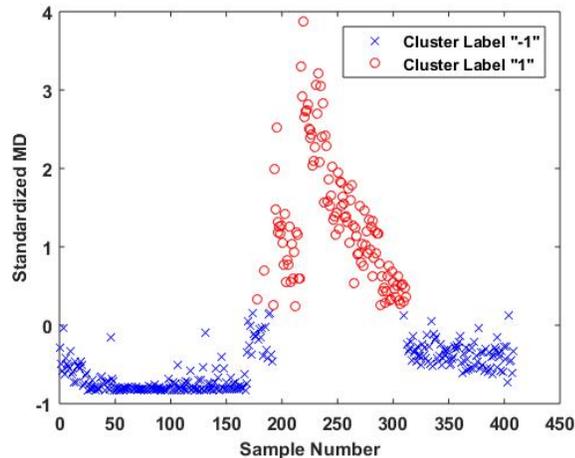


Figure 11. Test data shown with labels predicted by the naïve classifier using threshold found using k-means clustering

approach was chosen because these types of sensors do not require data acquisition equipment capable of high sample rates, and they are typically already installed on older hydro-turbines. Another benefit of this method is that it can be applied to data collected from other types of sensors including accelerometers, acoustic emission sensors, or pressure transducers, i.e., any sensor type that can be used to create a cavitation sensitivity parameter (CSP) sensitive to erosive cavitation.

The cavitation detection process presented here addresses issues unique to long term data collection by establishing a cavitation threshold from hydro-turbine ramp-down data and demonstrates how the process can be automated using an unsupervised learning algorithm. This strategy allows the thresholds to adapt to changes in running condition with minimal disruption to power production, and without human intervention. Here thresholds were established using a 90 second ramp-down, while cavitation surveys traditionally used to collect data for cavitation detection require stepping the hydro-turbine through a series of running conditions that can take several hours or even days to perform, and many more hours of manual analysis.

The method presented in this paper is a good starting point for researchers and hydro-turbine operators to better understand how to collect cavitation intensity data on a hydro-turbines for an extended period of time. The method can be used to identify a CSP, automate the training and classification process, and keep thresholds relevant through changes in operating conditions.

6. CONCLUSION

This paper presents both a novel method for creating adaptive cavitation thresholds as well as a machine learning frame-

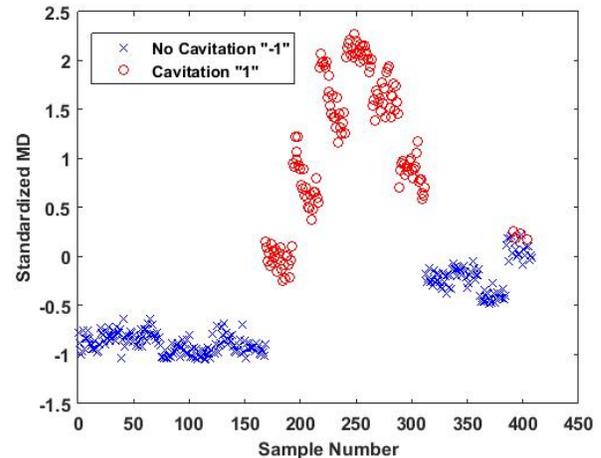


Figure 12. Test data with labels predicted by linear SVM model trained from data labeled using manually selected, single variable threshold and the training set PP3-CSP1, PP4-CSP1, PP1-CSP3, PP4-CSP3, PP3-CSP2

work for automated cavitation detection for hydro-turbines. Adaptive thresholds can be used to address issues encountered during long term cavitation detection caused by variability in the hydro-turbine's operating conditions – a critical part of collecting consistent intensity data for estimating erosion rates on hydro-turbine runners. The framework outlined here for automated cavitation detection provides a guideline for making data collection more practical and accessible for hydro-turbine operators and researcher wishing to estimate cavitation erosion rates and runner remaining useful life (RUL).

Adaptive cavitation thresholds are generated by first collecting sensor data from a hydro-turbine ramp-down, creating cavitation sensitivity parameter (CSP)s from the data and calculating the Mahalanobis distance (MD) to create clear separation between the healthy running state and conditions where the hydro-turbine is experiencing cavitation. This approach allows a new cavitation threshold to be generated quickly while minimizing impact on power production of the hydro-turbine, and being adaptable to variations in the turbine's running conditions. To automate the cavitation detection process, the cavitation threshold is used to create class labels for the ramp-down data that are used to train a supervised learning algorithm for classifying cavitation. Although domain knowledge is still required to select appropriate CSPs, the remainder of the process is automated by applying unsupervised learning to label the training set.

The results presented here show that K-Means and support vector machine (SVM)s for cavitation detection performed better than a process based on manually selected thresholds, demonstrating the usefulness of the machine learning framework. This approach provides hydro-turbine operators and re-

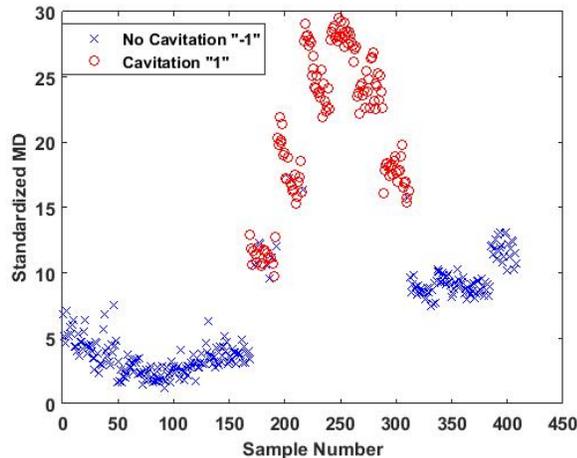


Figure 13. Test data with labels predicted by a 5th order polynomial SVM model trained from data labeled using multivariate threshold found with k-means clustering

searchers with a clear and effective way to perform automated cavitation detection and provides the basis for determining RUL.

7. FUTURE WORK

One important next step is to verify the methods for cavitation detection and intensity measurements by means of a long term study using a production hydroturbine. The larger data sets collected from such a study could be used to verify the accuracy and adaptability of the process demonstrated here and would lead to sufficient results to start correlating cavitation intensity measurements with erosion damage rates on turbine runners.

ACKNOWLEDGEMENT

The information, data, or work presented herein was funded in part by the Office of Energy Efficiency and Renewable Energy (EERE), U.S. Department of Energy, under Award Number DE-EE0002668 and the Hydro Research Foundation.

The authors wish to acknowledge the contributions of John Germann and James DeHaan for collecting the cavitation survey data and their guidance with analysis of the data. The authors further wish to acknowledge the code development assistance of Logan Schuelke.

The information, data or work presented herein was funded in part by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes and warranty, express or implied, or assumes and legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference

herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ACRONYMS

CSP	cavitation sensitivity parameter
DC	direct current
DFT	discrete Fourier transform
FFT	fast Fourier transform
MD	Mahalanobis distance
MW	Megawatts
RMS	root mean square
RUL	remaining useful life
SVM	support vector machine

REFERENCES

- An, D., Kim, N. H., & Choi, J.-H. (2013). Options for Prognostics Methods : A review of data-driven and physics-based prognostics. *Annual Conference of the Prognostics and Health Management Society*, 1–14. doi: 10.2514/6.2013-1940
- Avellan, F. (2004). Introduction to cavitation in hydraulic machinery. *The International Conference on Hydraulic Machinery ...*, 11–22. Retrieved from http://mmut.mec.upt.ro/mh/Conferinta{_}MH/102Avellan.pdf
- Bajic, B. (2002). Multidimensional Diagnostics of Turbine Cavitation. *Journal of Fluids Engineering*, 124(4), 943. doi: 10.1115/1.1511162
- Bajic, B., Services, K. C., GmbH, K., & Zithe, S. (2003). Methods for vibro-acoustic diagnostics of turbine cavitation Méthodes pour le diagnostic vibro-acoustique de la cavitation de turbine. *Analysis*, 41(1), 87–96.
- Bourdon, P., Farhat, M., Mossoba, Y., & Lavigne, P. (1999). Hydro Turbine Profitability and Cavitation Erosion. *Waterpower'99*, 1–10. Retrieved from [http://dx.doi.org/10.1061/40440\(1999\)76\\$\backslash\\$backslash\\$http://ascelibrary.org/doi/abs/10.1061/40440\(1999\)76](http://dx.doi.org/10.1061/40440(1999)76\backslashbackslash$http://ascelibrary.org/doi/abs/10.1061/40440(1999)76) doi: 10.1061/40440(1999)76
- Bourdon, P., Farhat, M., Simoneau, R., Pereira, F., Dupont, P., Avellan, F., & Dorey, J. (1996). Cavitation Erosion Prediction on Francis Turbines - Part 1 Measurements on the Prototype. In *Hydraulic machinery and cavitation: Proceedings of the xciii iahr symposium on hydraulic machinery and cavitation* (pp. 534–543).
- Bourdon, P., Simoneau, R., & Lavigne, P. (1989). A vibratory Approach to the Detection of Erosive Cavitation. In *In-*

- ternational symposium on cavitation noise and erosion in fluid systems*. ASME.
- Cencić, T., Hocevar, M., & Širok, B. (2014). Study of Erosive Cavitation Detection in Pump Mode of PumpStorage Hydropower Plant Prototype. *ASME J. Fluids Eng.*, 136(5), 51301. doi: 10.1115/1.4026476
- Cooley, J. W., & Tukey, J. W. (1964). An Algorithm for the Machine Calculation of Complex Fourier Series Author (s): James W . Cooley and John W . Tukey Published by : American Mathematical Society Stable URL : <http://www.jstor.org/stable/2003354> REFERENCES Linked references are available o. *Mathematics of Computation*, 19(90), 297–301.
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine learning*, 20(3), 273–297. doi: 10.1007/BF00994018
- De Maesschalck, R., Jouan-Rimbaud, D., & Massart, D. L. (2000). The Mahalanobis distance. *Chemometrics and Intelligent Laboratory Systems*, 50(1), 1–18. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0169743999000477> backslash\$nh<http://linkinghub.elsevier.com/retrieve/pii/S0169743999000477> doi: 10.1016/S0169-7439(99)00047-7
- Dorey, J.M.; Laperrousz, E.; Avellan, F.; Dupont, P.; Simoneau, R.; Bourdon, P. (1996). Cavitation Erosion Prediction on Francis Turbines - Part 3 Methodologies of Prediction. In *Hydraulic machinery and cavitation: Proceedings of the xciii iahr symposium on hydraulic machinery and cavitation*.
- Dorji, U., & Ghomashchi, R. (2014). Hydro turbine failure mechanisms: An overview. *Engineering Failure Analysis*, 44, 136–147. Retrieved from <http://dx.doi.org/10.1016/j.engfailanal.2014.04.013> doi: 10.1016/j.engfailanal.2014.04.013
- Dular, M., & Petkovšek, M. (2015). On the mechanisms of cavitation erosion Coupling high speed videos to damage patterns. *Experimental Thermal and Fluid Science*, 68, 359–370. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0894177715001508> doi: 10.1016/j.expthermflusci.2015.06.001
- Dular, M., Stoffel, B., & Širok, B. (2006). Development of a cavitation erosion model. *Wear*, 261(5-6), 642–655. doi: 10.1016/j.wear.2006.01.020
- Egusquiza, E., Valero, C., Valentin, D., Presas, A., & Rodriguez, C. G. (2015). Condition monitoring of pump-turbines. New challenges. *Measurement*, 67, 151–163. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0263224115000226> doi: 10.1016/j.measurement.2015.01.004
- Escaler, X., & Egusquiza, E. (2003). Vibration Cavitation Detection Using Onboard Measurements. *Symposium A Quarterly Journal In Modern Foreign Literatures*, 1–7.
- Escaler, X., Egusquiza, E., Farhat, M., Avellan, F., & Cousirat, M. (2006). Detection of Cavitation in Hydraulic Turbines. *Mechanical Systems and Signal Processing*, 983 – 1007.
- Escaler, X., Ekanger, J. V., Francke, H. H., Kjeldsen, M., & Nielsen, T. K. (2014). Detection of Draft Tube Surge and Erosive Blade Cavitation in a Full-Scale Francis Turbine. *Journal of Fluids Engineering*, 137(1), 011103. Retrieved from <http://fluidsengineering.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4027541> doi: 10.1115/1.4027541
- Figueiredo, E., Park, G., Farinholt, K. M., Farrar, C. R., & Lee, J.-R. (2012). Use of Time-Series Predictive Models for Piezoelectric Active-Sensing in Structural Health Monitoring Applications. *Journal of Vibration and Acoustics*, 134(4), 041014. Retrieved from <http://www.scopus.com/inward/record.url?eid=2-s2.0-84863806522{\&}partnerID=tZotx3y1> doi: 10.1115/1.4006410
- Flageul, R. F. P., a Archer, & C. (2012). Numerical and experimental investigations on cavitation erosion. *IOP Conference Series: Earth and Environmental Science*, 15(2), 22013. Retrieved from <http://stacks.iop.org/1755-1315/15/i=2/a=022013> doi: 10.1088/1755-1315/15/2/022013
- Francois, L. (2012). Vibratory detection system of Cavitation Erosion: Historic and Algorithm Validation. In *Proceedings of the eighth international symposium on cavitation* (pp. 325 –330).
- Gordon, J. L. (2001). Hydraulic turbine efficiency. *Canadian Journal of Civil Engineering*, 28(2), 238–253. doi: 10.1139/100-102
- Gregg, S. W., Steele, J. P. H., & Bossuyt, D. L. V. (2016). Feature Selection for Monitoring Erosive Cavitation on a Hydroturbine. *International Journal of Prognostics and Health Management, In Review*.
- Hammitt, F. G., & De, M. K. (1979). Cavitation damage prediction. *Wear*, 52(2), 243–262. doi: 10.1016/0043-1648(79)90066-8
- Hartigan, J. A. (1975). *Clustering Algorithms*.
- Hartigan, J. A., & Wong, M. A. (1979). A K-Means Clustering Algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1), 100–108. Retrieved from <http://www.jstor.org/stable/2346830>

- Heng, A., Zhang, S., Tan, A. C., & Mathew, J. (2009). Rotating machinery prognostics: State of the art, challenges and opportunities. *Mechanical Systems and Signal Processing*, 23(3), 724–739. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0888327008001489> doi: 10.1016/j.ymssp.2008.06.009
- Inacio, M., Lemos, A., & Caminhas, W. (2014). Fault Diagnosis with Evolving Fuzzy Classifier Based on Clustering Algorithm and Drift Detection. In *Annual conference of the prognostics and health management society* (Vol. 2014).
- ISO. (2005). *Mechanical Vibration – ISO 7919-5:2005: Evaluation of machine vibration by measurements on rotating shafts – Part 5: Machine sets in hydraulic power generating and pumping plants*.
- Jardine, A. K., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7), 1483–1510. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0888327005001512> doi: 10.1016/j.ymssp.2005.09.012
- Jian, W., Petkovšek, M., Houlin, L., Širok, B., & Dular, M. (2015). Combined Numerical and Experimental Investigation of the Cavitation Erosion Process. *Journal of Fluids Engineering*, 137(5), 051302. Retrieved from <http://fluidsengineering.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4029533> doi: 10.1115/1.4029533
- Kan, M. S., Tan, A. C., & Mathew, J. (2015). A review on prognostic techniques for non-stationary and non-linear rotating systems. *Mechanical Systems and Signal Processing*, 62-63, 1–20. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0888327015000898> doi: 10.1016/j.ymssp.2015.02.016
- Keogh, E., & Kasetty, S. (2002). On the need for time series data mining benchmarks. *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '02*, 102. Retrieved from <http://dl.acm.org/citation.cfm?id=775047.775062> doi: 10.1145/775047.775062
- Khelf, I., Laouar, L., Bouchelaghem, A. M., Rémond, D., & Saad, S. (2013). Adaptive fault diagnosis in rotating machines using indicators selection. *Mechanical Systems and Signal Processing*, 40(2), 452–468. Retrieved from <http://dx.doi.org/10.1016/j.ymssp.2013.05.025> doi: 10.1016/j.ymssp.2013.05.025
- Khurana, S., Navtej, & Singh, H. (2012). Effect of cavitation on hydraulic turbines- A review. *International Journal of Current Engineering and Technology*, 2(1), 172–177.
- The Knowledge Stream - Detecting Cavitation to Protect and Maintain Hydraulic Turbines. (2014). (Summer 2014). Retrieved from <https://www.usbr.gov/research/docs/updates/2014-14-cavitation.pdf>
- Kumar, P., & Saini, R. P. (2010). Study of cavitation in hydro turbines-A review. *Renewable and Sustainable Energy Reviews*, 14(1), 374–383. doi: 10.1016/j.rser.2009.07.024
- Lloyd, S. P. (1982). Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2), 129–137. doi: 10.1109/TIT.1982.1056489
- McKee, K. K., Forbes, G. L., Mazhar, I., Entwistle, R., Hodkiewicz, M., & Howard, I. (2015). A vibration cavitation sensitivity parameter based on spectral and statistical methods. *Expert Systems with Applications*, 42(1), 67–78. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0957417414004357> doi: 10.1016/j.eswa.2014.07.029
- Milligan, G. W., & Cooper, M. C. (1988). A study of standardization of variables in cluster analysis. *Journal of Classification*, 5(2), 181–204. doi: 10.1007/BF01897163
- Montgomery, D. C., & Runger, G. C. (2007). *Applied Statistics and Probability for Engineers* (4th ed.). John Wiley & Sons.
- Nandi, A. K., Liu, C., & Wong, M. L. D. (n.d.). *Intelligent Vibration Signal Processing for Condition Monitoring*, 1–15.
- Pal, N. R., & Pal, S. K. (1993). A review on image segmentation techniques. *Pattern Recognition*, 26(9), 1277–1294. doi: 10.1016/0031-3203(93)90135-J
- Pennacchi, P., Borghesani, P., & Chatterton, S. (2015). A cyclostationary multi-domain analysis of fluid instability in Kaplan turbines. *Mechanical Systems and Signal Processing*, 60-61, 375–390. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0888327015000163> doi: 10.1016/j.ymssp.2014.08.026
- Pollard, D. (1981). Strong Consistency of K-Means Clustering. *The Annals of Statistics*, 9(1), 135–140. doi: 10.1214/aos/1176345339
- Ramasso, E., & Saxena, A. (2014). Performance Benchmarking and Analysis of Prognostic Methods for CMAPSS Datasets. *International Journal of Prognostics and Health Management*(ISSN2153-2648), 1–15.
- Rus, T., Dular, M., Sirok, B., Hocevar, M., & Kern, I. (2007). An Investigation of the Relationship Between Acoustic Emission, Vibration, Noise, and Cavitation Structures on a Kaplan Turbine. *Journal of Fluids Engineering*,

129(September), 1112.

- Samanta, B., Al-Balushi, K. R., & Al-Araimi, S. A. (2003). Artificial neural networks and support vector machines with genetic algorithm for bearing fault detection. *Engineering Applications of Artificial Intelligence*, 16(7-8), 657–665. doi: 10.1016/j.engappai.2003.09.006
- Saxena, A., Celaya, J. R., Saha, B., Saha, S., & Goebel, K. (2009). On Applying the Prognostic Performance Metrics. *Proceedings of the annual conference of the prognostics and health management society*, 1–16.
- Schmidt, H., Kirschner, O., Riedelbauch, S., Necker, J., Kopf, E., Rieg, M., ... Mayrhuber, J. (2014). Influence of the vibro-acoustic sensor position on cavitation detection in a Kaplan turbine. *IOP Conference Series: Earth and Environmental Science*, 22(5), 052006. Retrieved from <http://iopscience.iop.org/1755-1315/22/5/052006><http://stacks.iop.org/1755-1315/22/i=5/a=052006?key=crossref.7e388f6f614fcf79fa07b58f190a2879> doi: 10.1088/1755-1315/22/5/052006
- Toyota, T., Niho, T., & Chen, P. (2000). Condition monitoring and diagnosis of rotating machinery by Gram-Charlier expansion of vibration signal. In *Knowledge-based intelligent engineering systems and allied technologies*.
- US Department of the Interior Bureau of Reclamation. (2014). *The Knowledge Stream - Hydropower and Renewable Energy Issue* (No. Summer). Retrieved from www.usbr.gov/research/docs/ks-2014-04.pdf
- U.S. Energy Information Administration. (2015). *Electric Power Monthly: with data for May 2015* (Tech. Rep. No. May).
- Varga, JJ and Sebestyen, Gy and Fay, A. (1969). Detection of cavitation by acoustic and vibration-measurement methods. *La houille blanche*(2), 137–150.
- Widodo, A., & Yang, B.-S. (2007). Support vector machine in machine condition monitoring and fault diagnosis.
- Witten, I. H., & Frank, E. (2005). *Data Mining: practical machine learning tools and techniques*. Morgan Kaufmann.
- Wolff, P. J., Jones, R. K., & March, P. (2005). *Evaluation of Results from Acoustic Emissions-Based Cavitation Monitor, Grand Coulee Unit G-24 Cavitation Monitoring System Comparison Tests, Grand Coulee Project Final Report* (Tech. Rep.).
- Wu, X., Kumar, V., Ross, Q. J., Ghosh, J., Yang, Q., Motoda, H., ... Steinberg, D. (2008). *Top 10 algorithms in data mining* (Vol. 14) (No. 1). doi: 10.1007/s10115-007-0114-2

APPENDIX

Classifier	Training and Testing Set	% Correct	I_{total}
naïve, with manually selected threshold	Proximity Probe 1	95.3	2457
	Proximity Probe 2	95.3	2473
	Proximity Probe 3	93.6	2696
	Proximity Probe 4	92.2	2395
naïve, with K-Means selected threshold	Proximity Probe 1	91.2	2433
	Proximity Probe 2	94.6	2473
	Proximity Probe 3	89.0	2661
	Proximity Probe 4	85.1	2395
SVM, trained with manually selected threshold	PP2-CSP1, PP2-CSP3, PP3-CSP3, PP2-CSP2	95.6	2446
	PP2-CSP1, PP2-CSP3, PP3-CSP3, PP2-CSP2, PP3-CSP2	95.3	2433
	PP1-CSP3, PP2-CSP3, PP3-CSP3, PP1-CSP2, PP2-CSP2	95.1	2520
SVM, trained with K-Means selected threshold	PP3-CSP1, PP4-CSP1, PP1-CSP3, PP4-CSP3, PP3-CSP2	99.0	2638
	PP3-CSP1, PP4-CSP1, PP1-CSP3, PP2-CSP3, PP4-CSP3, PP4-CSP2	98.7	2633
	PP3-CSP1, PP4-CSP1, PP1-CSP3, PP2-CSP3, PP4-CSP3, PP1-CSP2, PP4-CSP2	98.7	2640
	PP3-CSP1, PP4-CSP1, PP1-CSP3	98.5	2625

Table 2. Classifier test results for single variable thresholds

Classifier	Training and Testing Set	% Correct	I_{total}
non-linear SVM trained with manually selected multivariate threshold	CSP2, CSP3, 1st order	94.1	2711
	CSP2, CSP3, 2nd order	94.1	2711
	CSP1, CSP2, CSP3, 3rd order	94.4	2720
	CSP1, CSP2, CSP3, 4th order	96.8	2829
	CSP1, CSP2, CSP3, 5th order	97.8	2866
	CSP1, CSP2, CSP3, 6th order	97.3	2833
	CSP1, CSP2, CSP3, 7th order	96.8	2798
	CSP1, CSP2, CSP3, 8th order	96.1	2740
non-linear SVM trained with K-means selected multivariate threshold	CSP2, CSP3, 1st order	94.1	2711
	CSP3, 2nd order	93.6	2777
	CSP3, 3rd order	94.0	2696
	CSP1, CSP2, CSP3, 4th order	95.1	2916
	CSP1, CSP2, CSP3, 5th order	97.5	2891
	CSP1, CSP2, CSP3, 6th order	97.3	2891
	CSP1, CSP2, CSP3, 7th order	96.8	2951
	CSP1, CSP2, CSP3, 8th order	96.1	2891
naïve, with multivariate, manually selected threshold	CSP3	94.1	2711
naïve, with multivariate, K-means selected threshold	CSP3	94.1	2711

Table 3. Classifier test results for multivariate thresholds