

## MULTI-STAGE FAULT WARNING FOR LARGE ELECTRIC GRIDS USING ANOMALY DETECTION AND MACHINE LEARNING

SANJEEV RAJA AND ERNEST FOKOUÉ

*Abstract.* In the monitoring of a complex electric grid, it is of paramount importance to provide operators with early warnings of anomalies detected on the network, along with a precise classification and diagnosis of the specific fault type. In this paper, we propose a novel multi-stage early warning system prototype for electric grid fault detection, classification, subgroup discovery, and visualization. In the first stage, a computationally efficient anomaly detection method based on quartiles detects the presence of a fault in real time. In the second stage, the fault is classified into one of nine pre-defined disaster scenarios. The time series data are first mapped to highly discriminative features by applying dimensionality reduction based on temporal autocorrelation. The features are then mapped through one of three classification techniques: support vector machine, random forest, and artificial neural network. Finally in the third stage, intra-class clustering based on dynamic time warping is used to characterize the fault with further granularity. Results on the Bonneville Power Administration electric grid data show that i) the proposed anomaly detector is both fast and accurate; ii) dimensionality reduction leads to dramatic improvement in classification accuracy and speed; iii) the random forest method offers the most accurate, consistent, and robust fault classification; and iv) time series within a given class naturally separate into five distinct clusters which correspond closely to the geographical distribution of electric grid buses.

### 1. INTRODUCTION

Electric grids are of vital importance in the infrastructure of modern society, serving to ensure the continuous supply of electricity to households and industries. Grid failures lead to significant financial losses for companies and inconvenience for consumers and maintenance personnel. Automatic monitoring of large complex electric grids has thus received considerable attention from the research community in recent years. Many authors have addressed this challenging problem from a variety of angles using a combination of tools from computer science, electrical engineering, statistics, machine learning and artificial intelligence. In this work we present a novel multistage system using statistical anomaly detection and machine learning techniques to detect, classify, and diagnose electrical faults. While our methodology is applicable to any general time series data, we use the Bonneville

---

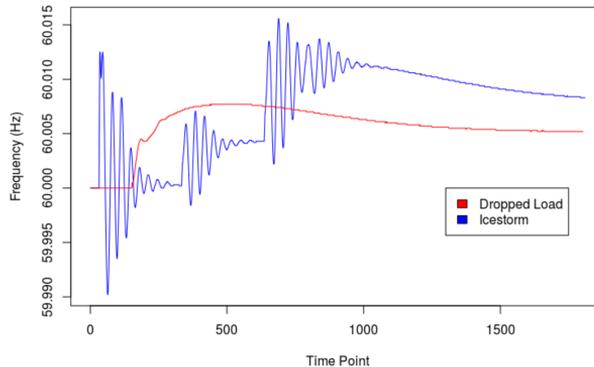
*MSC (2010):* primary 62H30; secondary 68T05, 68T10.

*Keywords:* electric grid, anomaly detection, time series, machine learning, classification, neural network, random forest, support vector machine, dimensionality reduction, autocorrelation, clustering, dynamic time warping.

Power Administration (BPA) dataset to develop and evaluate our techniques. The BPA electric grid comprises 126 distinct stations or buses, and spans the states of Washington and Oregon. Our dataset is represented as a collection of 126 time series, one for each station, and each time series contains 1800 observations corresponding to 60 seconds of data at a sampling rate of 30 Hz. Three sets of observations are available, namely frequency, voltage, and phase angle metrics.

Stations in a power grid can fail in several distinct ways. Each failure scenario is characterized by unique time-varying frequency, voltage, and phase angle profiles, and entails specialized repair by operators. We will refer to these scenarios as fault classes. In the BPA dataset, each time series represents the occurrence of exactly one fault over the specified time period. The dataset contains 16 labeled fault classes corresponding to commonly encountered failures. We possess 126 instances of each fault class corresponding to the 126 stations in the BPA grid, for a total of approximately 2000 data instances. Of the 16 fault classes, several were observed to exhibit near identical temporal behavior. To simplify our classification models and avoid overfitting, highly overlapping fault classes were merged. The criteria for fault merging was visual inspection of the frequency, voltage, and phase angle time series. The result is the following nine fault class labels: *Dropped Load*, *Open AC*, *Open DC*, *Open Generator*, *GMD 2*, *Ice Storm*, *McNary Attack*, *Ponderosa*, and *Quake 1*. Our goal is to automatically identify and communicate the presence and type of fault in real time so that preemptive action can be taken.

Figure 1 depicts typical frequency time series for the *Ice Storm* and *Dropped Load* fault classes. The temporal characteristics for these fault types are visibly different.



**Figure 1.** Frequency time series of *Ice Storm* and *Dropped Load* faults.

Recent literature on fault detection and analysis in electric grid systems has largely focused on machine learning approaches [1–8]. Notably, in [5] a control system approach involving micro-controllers and other electrical equipment was introduced to monitor faults based solely on current fluctuations. In [6] and [7], a fault detection method was combined with classification by an Artificial Neural Network. In [8], a Long-Short Term Memory (LSTM) neural network architecture was proposed to perform fault classification. To our knowledge, investigations

exploring random forests or support vector machines, models which are computationally efficient and provide high classification accuracy, are relatively sparse in the literature. Dimensionality reduction and feature selection techniques to reduce computational cost are also not well-developed. Furthermore, subgroup discovery within a fault class using unsupervised learning has not been attempted. Our work is thus novel in that it combines traditional statistical methods such as anomaly detection with advanced machine learning, dimensionality reduction, and visualization, into a cohesive system for real-time fault diagnosis and prediction.

Prior work by the authors on BPA-specific electric grid data has focused almost exclusively on visualization and exploratory analysis [9]. This included identifying the Gaussian nature of the distribution of electric frequency measurements across the grid at a specific time point, and generating heat maps of the rate of change of grid parameters across stations. All of the aforementioned analysis was performed for one specific class. In this work, we progress further by analyzing data across multiple classes to develop a warning system that is both highly accurate and computationally efficient, rendering it feasible for expedited implementation at a grid scale.

The rest of this paper is organized as follows: Section 2 provides an overview of the proposed multi-stage fault warning system. Section 3 describes the first stage of the system, namely dynamic anomaly detection. Section 4 presents dimensionality reduction and three fault classification methods. Section 5 describes intra-class clustering. Experimental results are presented in Section 6, and concluding remarks are given in Section 7.

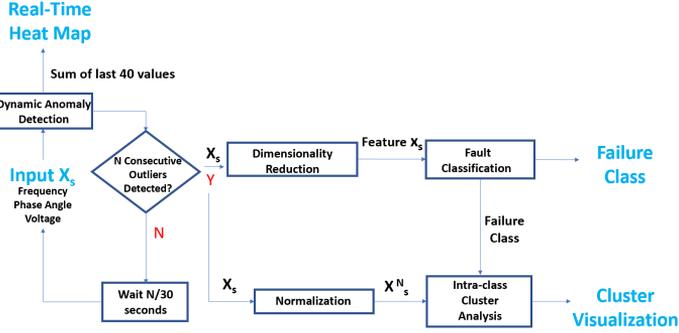
## 2. OVERVIEW OF EARLY WARNING SYSTEM

We present a framework on which to build a real-time early warning system for BPA electric grid failures, depicted in Figure 2. The 126 BPA electric grid stations each produce observations for frequency, voltage, and phase angle. We represent the time series for a station  $s$  from time  $t = 1$  to  $t = T$  as a vector  $X_s = (X_{s,1}, \dots, X_{s,t}, \dots, X_{s,T})$ , where each  $X_{s,t}$  can be a frequency, voltage, or phase angle measurement. The  $T$ -dimensional vectors  $X_s$  are the inputs to the early warning system.

In Stage 1, we perform dynamic anomaly detection on the time series. The results, along with those of the other buses, can be visualized on a heat map. If  $n$  consecutive anomalous time points (i.e. outliers) are detected in  $X_s$ , where  $n$  is an experimentally determined threshold, we assume a fault has occurred, freeze further input to the system from that station, and proceed to Stage 2. Otherwise, we wait  $\frac{n}{30}$  seconds (corresponding to a sampling frequency of 30 Hz), receive the latest time series, and conduct Stage 1 analysis again.

In Stage 2, we perform dimensionality reduction on  $X_s$  to produce a compact representation  $x_s$  of the time series. We then pass this representation into a classifier (trained offline) to determine the predicted fault class of the time series from a total of 9 possible classes.

In Stage 3, having identified the fault class of  $x_s$ , we perform clustering of data within each class with the aim of gathering insight on the source and severity of a given fault type.



**Figure 2.** Diagram of proposed early warning system. The main components of the process are outlier detection, dimensionality reduction, fault classification, and intra-class clustering.

Each of the stages are described in more detail in the following sections.

### 3. DYNAMIC ANOMALY DETECTION

We employ a simple, robust, and computationally efficient criterion for detecting the presence of anomalous time points or outliers in the electric grid time series data. For a given station  $s$ , we map its frequency, voltage, or phase angle time series  $X_s$  to an outlier vector  $O_s$  of the same length. For each time point  $X_{s,t}$  in  $X_s$ , we declare the point a moderate outlier if the following condition holds:

$$X_{s,t} > Q_3 + 1.5 \times \text{IQR} \quad \text{or} \quad X_{s,t} < Q_1 - 1.5 \times \text{IQR}$$

where  $Q_1$  and  $Q_3$  are the first and third quartile of  $X_{s,1,\dots,t}$  respectively, and  $\text{IQR}$  is the interquartile range of  $X_{s,1,\dots,t}$ . If this condition is true,  $O_{s,t}$  is set to 1. If the following condition holds:

$$X_{s,t} > Q_3 + 3 \times \text{IQR} \quad \text{or} \quad X_{s,t} < Q_1 - 3 \times \text{IQR}$$

the time point is considered a severe outlier and the value of  $O_{s,t}$  is set to 2. If neither of the above two conditions are true, the point is not considered an outlier, and  $O_{s,t}$  is set to 0.

We thus generate in real time a ternary outlier vector  $O_{s,t}$  describing anomaly state as the frequency, voltage, or phase angle measurements from a bus are received. Only time series that exhibit a number of consecutive anomalous observations are passed to the fault classification stage.

### 4. FAULT CLASSIFICATION

We use supervised learning techniques to classify the electric grid time series into one of the nine labeled fault classes. As a first step we extract compact features from the raw electrical time series via a dimensionality reduction step, described next.

#### 4.1. Dimensionality reduction

A crucial aspect of any classification task is identifying discriminative features of the data on which to perform classification. Raw data often disguises these features and contains considerable redundancy, reducing the accuracy and efficiency of classification. We opt to use the autocorrelation (ACF) function to produce a time-invariant, compact feature representation of the data. ACF returns a vector of values between  $-1$  and  $1$  representing the correlation of a time series with lagged copies of itself. Note that the ACF output has constant, specified dimensionality regardless of the dimensionality of the input.

Since ACF is only meaningful on stationary time series, we conduct first-order differencing of the time series  $X_s$  to ensure stationarity. The differenced time series  $\hat{X}_s$  is given by:

$$\hat{X}_{s,t} = X_{s,t} - X_{s,t-1}, \quad \text{for } t = 2, \dots, T, \quad \text{with } \hat{X}_{s,1} = 0.$$

The autocovariance function at lag  $h$  is:

$$\begin{aligned} \gamma_s(h) &= \text{cov}(\hat{X}_{s,t+h}, \hat{X}_{s,t}) \\ &= \mathbb{E} \left[ (\hat{X}_{s,t+h} - \mu_s)(\hat{X}_{s,t} - \mu_s) \right]. \end{aligned}$$

The corresponding autocorrelation function is:

$$\rho_s(h) = \frac{\gamma_s(h)}{\gamma_s(0)} = \text{corr}(\hat{X}_{s,t+h}, \hat{X}_{s,t}).$$

For electrical data from a given substation, and for a fixed lag  $h$ , we obtain a scalar autocorrelation value. By computing this value for  $K$  different lags, we obtain a  $K$ -dimensional feature vector denoted  $x_s$  which is then used as input to the classifiers described next.

#### 4.2. Classification techniques

The input to the classifier is feature  $x_s \in \mathbb{R}^K$ , and the output is a label  $y_s$  corresponding to one of nine fault classes. We applied three well known machine learning methods to classify BPA data. We now briefly present these techniques, deferring the reader to the respective references for more detailed descriptions.

**4.2.1. Support vector machine.** A Support Vector Machine (SVM) is a supervised learning method for data classification tasks [10]. Given a set of training data with class labels, SVM learns hyperplanes (i.e. linear decision boundaries) that optimally divide the data by class. A test point is then assigned to a class based on which side of the hyperplane it falls into. SVM handles more complex classification problems with nonlinear decision surfaces by first transforming input vectors through a nonlinear mapping to a very high-dimensional feature space prior and constructing linear hyperplanes in this space. The derivation of the hyperplane is outside the scope of this paper (see [10]); we simply note that the parameters of the hyperplane depend only on a few support vectors,  $x_{sj}$  which are in effect the data points closest to the decision boundary. SVM is well-suited to classification of high-dimensional, sparse data, and is fast to execute, making it an apt technique for classifying BPA electric grid data.

For illustration consider a 2-class problem with inputs  $x_s$  and output class labels  $+1$  and  $-1$ . The SVM classifier is expressed by the following equation:

$$\hat{f}(\mathbf{x}_s) = \text{sign} \left( \sum_{j=1}^N \hat{\alpha}_{s_j} y_{s_j} K(\mathbf{x}_{s_j}, \mathbf{x}_s) + \hat{b} \right)$$

where  $x_{s_j}$  are  $N$  support vectors,  $K$  is a kernel function that takes the input samples into a high-dimensional space, and  $\hat{\alpha}_{s_j}$  and  $\hat{b}$  are the parameters of the separating hyperplane learned during training. The binary classification problem can be readily extended to multiple classes [10].

**4.2.2. Random forest.** Random Forests (RF) enable a probabilistic ensemble learning method for data classification [11]. The basic block of an RF is a decision tree which recursively splits the  $K$ -dimensional feature space until a partition of  $P$  classes is produced (in our application,  $P = 9$ ). In the case of a binary decision tree, each node of the tree splits the  $K$ -dimensional space into two partitions. Repeated splits are performed until the  $P$ -sized partition is achieved. The parameters of the splits are optimized during a training phase to minimize an overall classification error. One of the shortcomings of decision trees is that they tend to overfit on training data. To mitigate this issue, the RF algorithm aggregates decisions from multiple decision trees. Essentially an RF is comprised a forest of trees, whereby each tree is learned from a random sample of the training data, and the optimal split at each node of a tree is chosen from a random sample of the features of the training data. The output of the RF classifier is the mode (i.e. majority vote) of the class labels predicted by the individual trees, along with a probability of the predicted class membership. Due to the feature diversity offered by a large collection of trees, RF classifiers are generally robust to overfitting, an important consideration for a task such as electrical fault classification with potentially complex decision boundaries.

**4.2.3. Artificial neural network.** An artificial neural network (ANN) is a group of interconnected nodes organized into layers: an input layer with one node for each input feature, several hidden layers, and an output layer with one node for each possible output [12]. Every inter-node connection in the network carries an associated weight and a function that maps an input to a known output.

The computations at the single node of Figure 3 are given by:

$$\eta(\mathbf{x}) = o_j = f_j(\beta_j + \sum_i w_{ji} x_i)$$

where  $\beta_j$  is a learnable scalar bias term. An ANN combines multiple nodes in multiple cascaded layers to realize arbitrarily complex classification and regression functions. Several choices exist for the activation function, such as the logistic sigmoid function:

$$\psi(z) = \frac{1}{1 + e^{-z}}$$

and hyperbolic tangent function:

$$\psi(z) = \tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}} = \frac{1 - e^{-2z}}{1 + e^{-2z}}.$$

We choose the hyperbolic tangent function due to its favorable convergence properties.

During training, the network weights  $w_{ji}$  are first initialized, usually with random coefficients. The network processes training samples one at a time, comparing the predicted output of the network to the ground truth. This comparison yields an error that is back-propagated through the system, and the node connection weights  $w_{ji}$  are adjusted to minimize training error. This process iterates until convergence (see [12] for a detailed overview of ANN training). ANN classifiers are highly tolerant to noisy data and can learn arbitrarily complex decision boundaries, an important criterion for multinomial classification. In our experiments, we employ a 3-layer ANN, whose details are given in Section 6.

## 5. INTRA-CLASS UNSUPERVISED LEARNING

Individual stations in an electric grid can respond differently to a given fault based on geography, size, staffing, and a number of other factors. We use unsupervised learning techniques, namely time series clustering, to discover meaningful subgroups of grid stations within a fault type. These subgroups could potentially correspond to varying severity of a fault or other differentiating factors. With intra-class analysis, grid operators gain access to more granular information such as specific geographical areas that have been affected most severely and require immediate intervention. While the fault classifier in the previous stage operates on a compact ACF representation  $x_s$ , in this stage we operate on the original raw time series  $X_s$  to uncover temporal structure within each class.

Prior to clustering, we normalize the time series so that all values are in the range of  $[0, 1]$ . Normalization augments the disparity between times series, allowing for greater clustering acuity. For each  $X_{s,t}$  in a time series  $X_s$ , the normalized value  $X_{s,t}^N$  is given by

$$X_{s,t}^N = \frac{X_{s,t} - \min X_s}{\max X_s - \min X_s}.$$

### 5.1. Dynamic time warping

Any clustering technique requires a metric that defines distance between samples. Since we are dealing with time series, we use dynamic time warping (DTW) [13] to compute the similarity (or distance) between two temporal sequences. To briefly review DTW, given two time series  $U$  and  $V$  of length  $m$  and  $n$ , respectively, an  $m \times n$  distance matrix  $D$  is constructed with elements  $D_{ij}$  representing the pairwise distance between points  $U_i$  and  $V_j$ . Euclidean distance is commonly used, such that

$$D_{ij} = \sqrt{(U_i - V_j)^2}.$$

A warping path  $w$  is defined as a contiguous sequence of  $k$  matrix elements that satisfies the following two conditions:

- (1) Boundary conditions:  $w_1 = (1, 1)$  and  $w_k = (m, n)$ ;
- (2) Continuity and monotonicity: if  $w_i = (a, b)$  then  $w_{i-1} = (a', b')$  where  $0 \leq a - a' \leq 1$  and  $0 \leq b - b' \leq 1$ .

The cost of the path is defined as the sum of the elements (distances) traversed by the path. DTW seeks the path with minimum cost, the latter being the DTW distance between the two sequences. DTW is often sped up by limiting the set of valid paths to a limited region of matrix  $D$  around the diagonal.

## 5.2. Partitioning around medoids

Given a distance metric, we next proceed to cluster the data using the Partitioning Around Medoids (PAM) algorithm [14]. PAM is a variant of the well-known K-means clustering algorithm [15]. However instead of representing each cluster with its centroid as is done in K-means, PAM represents each cluster with an exemplar data point, referred to as the medoid (or “middle point”). More crucially, PAM admits more general distance metrics, in contrast to K-means which uses only squared Euclidean distance. Hence PAM is favorable for our application where DTW defines distances between time series.

With the task of partitioning the data in a given fault class into  $L$  clusters, PAM proceeds as follows:

- (1) Select  $L$  out of the data points as initial medoids;
- (2) Associate each data point with the closest medoid, with distance measured by DTW;
- (3) Compute the total cost, which is the sum of distances of points to their assigned medoids;
- (4) While the total cost decreases:
  - (a) For each medoid  $m$  and for each non-medoid sample  $d$ 
    - (i) Swap  $m$  and  $d$ , reassign all points to the closest medoid and compute total cost;
    - (ii) If the total cost increased in the previous step, undo the swap.

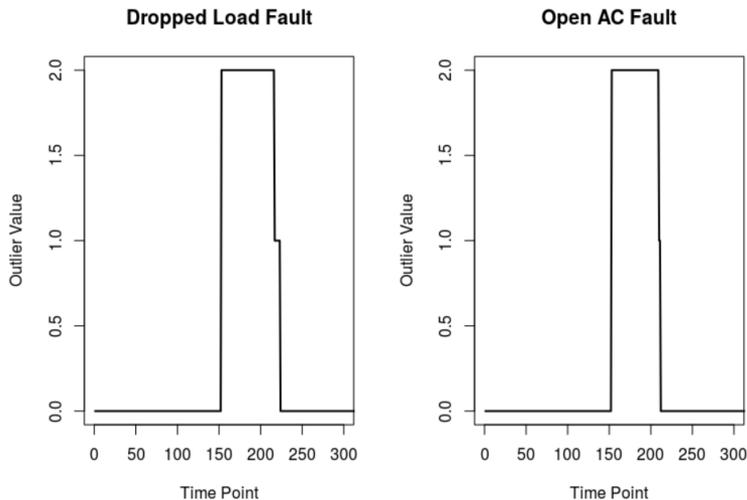
## 6. EXPERIMENTAL RESULTS

We now present results for the BPA data. We used the  $R$  statistical language [16] to develop all the analysis. Recall that the data contains 126 time series  $X_s$ , one for each station in the BPA grid. Most time series contain 1802 time samples, corresponding to just over 60 seconds of frequency, voltage, or phase angle data at a sampling rate of 30 Hz, while a few time series are longer, containing up to 3000 samples. In preliminary experiments, we found frequency observations to be more discriminative for fault classification than voltage and phase angle; thus frequency data was used for all experiments. Any further mention of time series refers specifically to frequency time series.

### 6.1. Anomaly detection

Applying our quartile criterion to detect outliers in the frequency data, we found that the onset of a fault is characterized by approximately 70 consecutive severe outliers followed by a return to normalcy due to the self-adjusting nature of the outlier detection criterion. This is depicted in Figure 3.

Based on this signature of fault onset, we propose setting our early warning system outlier threshold at 70. In other words, when we detect 70 consecutive



**Figure 3.** Outlier vector for *Dropped Load* and *Open AC* faults generated using quartile criterion. Faults are reliably characterized by a sudden spike of 70 anomalous time points.

outliers relative to the accepted 60 Hz frequency value, we assume a fault has occurred and proceed to the second stage for more detailed classification.

We also used the outlier vector to generate an intuitive visualization of the BPA electric grid, shown in Figure 4. The colors of the scatter plot correspond to the sum of the last 40 values of the outlier vector  $O_s$  for a given station (ranging from 0-80). This provides unique insight into the location and severity of the fault, and can be updated in real-time to provide operators a live picture of the health of the grid.

## 6.2. Dimensionality reduction

We conducted first-order differencing of the frequency series and applied the ACF function to obtain a compact representation. The maximum number of lags we used was 20. ACF thus condenses our raw data  $X_s$  containing 1800 to 3000 time points to a time-invariant 20-dimensional feature vector  $x_s$ . Figure 5 shows the ACF feature vector of the differenced frequency series for the *Icestorm* and *Dropped Load* faults.

We performed an ablation study to determine the accuracy of SVM classification with and without dimensionality reduction. Using the original frequency data with 1802 dimensions, it took 21.2 seconds to train an SVM classifier on 1602 examples, and 2.29 seconds to predict the fault class of 399 test time series, with 53.8 percent of test samples being labelled correctly. Using the 20-dimensional ACF representation on the same training and test data, training took 0.25 seconds, testing took 0.031 seconds, and there was a dramatic increase in classification accuracy to 96.0 percent. Clearly, ACF dimensionality reduction results in a highly discriminative, compact representation, allowing for real-time, accurate classification that cannot be accomplished using raw data. We also experimented with the

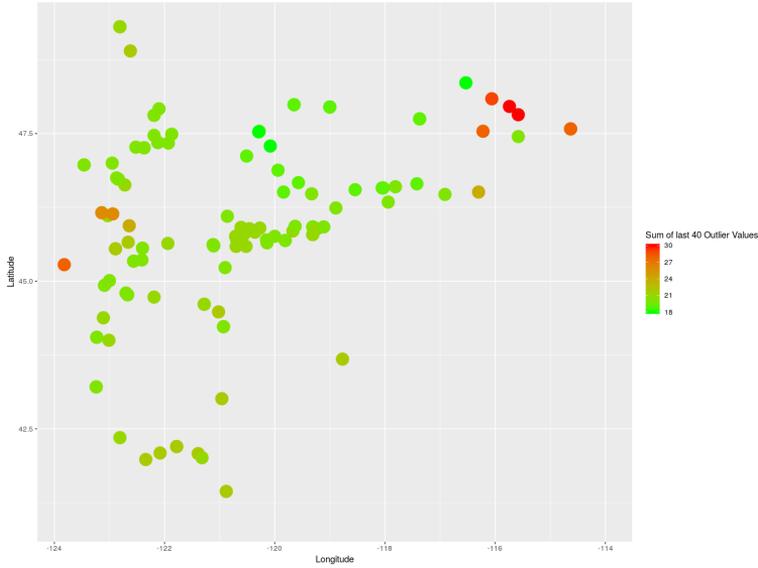


Figure 4. Severity of *Icestorm* fault across BPA electric grid stations.

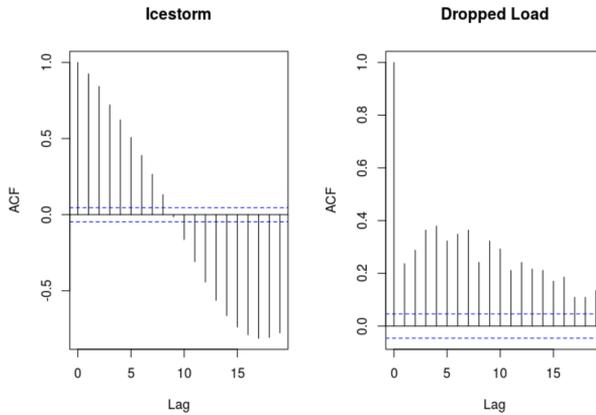


Figure 5. Autocorrelation feature vector for differenced frequency series corresponding to *Icestorm* and *Dropped Load* faults.

partial autocorrelation function [17] and spectral periodogram [18] obtained via a Fast Fourier Transform and found that both methods were inferior to the ACF in terms of discriminability, classification accuracy, and robustness to inputs of varying dimensions.

### 6.3. Fault classification

We applied the three classification models introduced above to the task of classifying the BPA time series into nine fault scenarios. The two fault classes that

were a result of consolidation contained approximately 500 data examples each, while the others contained 126 examples each corresponding to the 126 buses in the BPA electric grid.

**6.3.1. Support vector machine.** We used the radial basis function with a gamma value of 0.05 as our SVM kernel. We set our soft margin cost parameter to 1. Out of a total of 2001 data examples, 80 percent (1602) were used to train the classifier, and the remaining 20 percent (399) were reserved for testing. The training input to the classifier was a set of 20-dimensional ACF vectors of first-order differenced time series along with corresponding fault class labels. At inference, the input to the classifier was an ACF vector, and the output was the predicted fault class. The mean classification accuracy of the SVM classifier on the testing data, calculated over 100 trials on different training-testing splits, was 97.2 percent.

**6.3.2. Random forest.** We implemented an RF classifier with 500 trees, determined to be optimal from hyperparameter tuning. The same protocols used for training and evaluating the SVM classifier were also used for the RF classifier. The mean classification accuracy of the RF classifier on the testing data, calculated over 100 trials on different training-testing splits, was 98.9 percent.

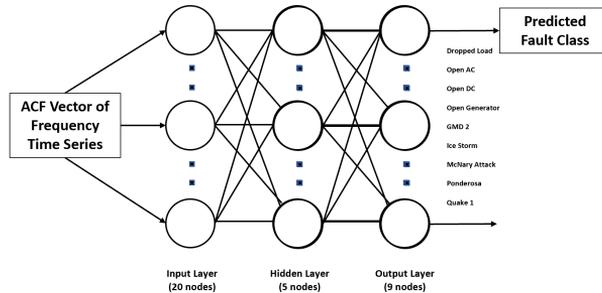
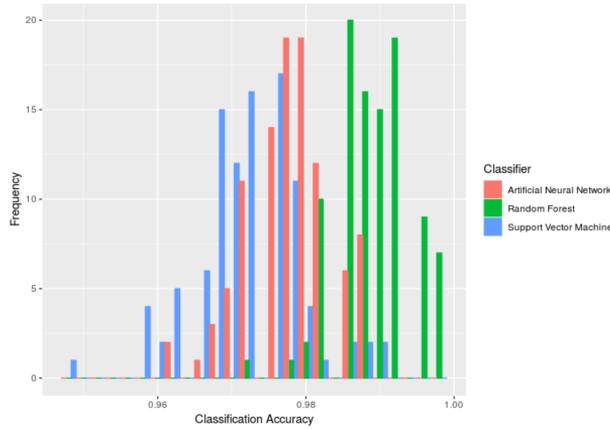


Figure 6. Artificial neural network for fault classification.

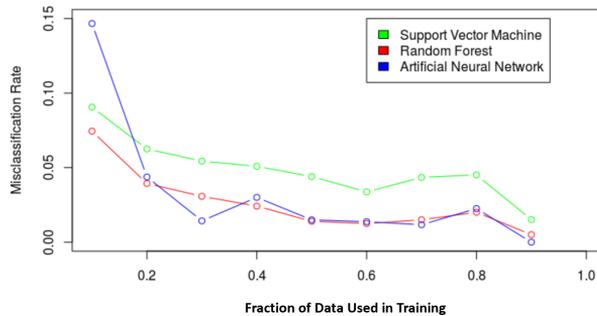
**6.3.3. Artificial neural network.** We implemented the ANN classifier depicted in Figure 6. The network consists of an input layer of 20 nodes corresponding to the 20-dimensional ACF vector, a single hidden layer of 5 nodes, and an output layer of 9 nodes for each target fault class. The same protocols used for the previous classifiers were used to train and evaluate the ANN. The mean classification accuracy calculated over 100 trials on different training-testing splits was 97.8 percent.

**6.3.4. Comparison of classifiers.** Figure 7 compares histograms of accuracy scores for the three classifiers. Each classifier was run 100 times, with each iteration using a unique train-test split stratified by fault type. We note that the classification accuracy of the RF classifier is consistently higher than that of SVM



**Figure 7.** Histogram of RF, SVM, and ANN classifier accuracy over 100 trials (accuracy=1 corresponds to 100 percent of test samples being classified correctly).

and ANN. RF classification accuracies are skewed left and are more tightly clustered around their mean of 98.9 percent. SVM classification accuracies are more nearly Gaussian around 97.2 percent with a greater spread, while ANN accuracies are skewed right but with a lower mean of 97.8 percent. We can infer that not only does RF on average outperform SVM and ANN on our classification task, but it also displays less variability and therefore greater reliability across many iterations of classification.

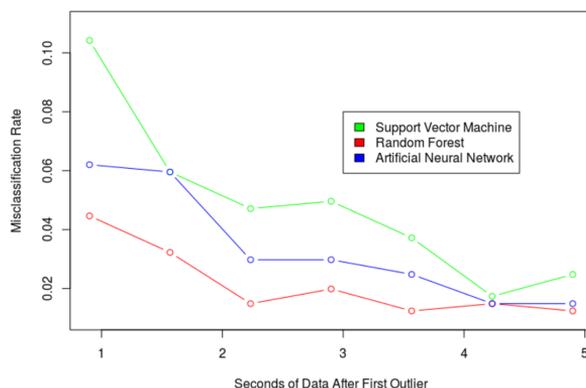


**Figure 8.** Classifier performance as a function of training budget.

Figure 8 depicts the performance of the 3 classifiers as a function of the size of the training set, represented as a fraction of the total number of data samples. It is clear that when at least 20 percent of the data (399 time series) were used in training, both ANN and RF consistently outperformed SVM, with the former two being similar in performance. However, when less than 20 percent of the data was used in the training set, ANN’s misclassification rate rose to nearly 15 percent. ANN’s considerable drop in performance can be attributed to its complexity as a model. ANN is able to learn arbitrarily complex decision boundaries, likely causing it to overfit in the presence of sparse training data. SVM and RF are

comparatively simpler models that appear more robust to overfitting on this data set. The discrepancy among classifier performances on sparse training sets is highly revealing. For the BPA dataset, both RF and ANN appear to be equally viable candidates for classification, as there are no data limitations. However, when considering the application of our early warning system to other electric grids or domains, availability of training data may be an important consideration. In sparse training scenarios, our results strongly suggest that RF is the premier candidate for fault classification.

**6.3.5. Preemptive classification.** Classification results reported thus far have been on complete time series containing 1802 time points, or approximately 60 seconds of data. Realistically, we aim to classify faults before they have run their entire course. To this end, we couple the outlier detection method with classification. Upon identification of 70 consecutive outliers, we aim to classify the time series as quickly as possible to provide ample time for operators to intervene. We trained the RF, SVM, and ANN classifiers on features derived from varying numbers of time samples captured immediately after the occurrence of the first outlier in the 70-outlier series. ACF was performed on the time series prior to classification. Figure 9 plots classifier performance as a function of the number of temporal samples used for feature definition.



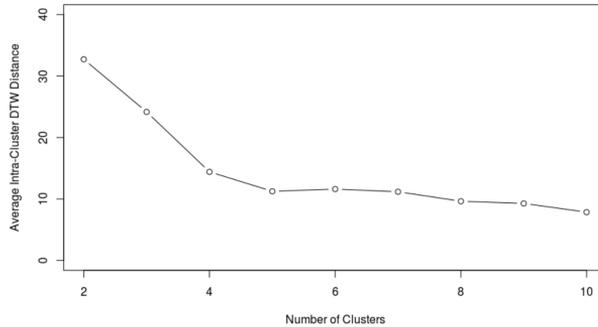
**Figure 9.** Misclassification Rates of RF, SVM, and ANN as a function of number of temporal samples used to define input features.

The results in Figure 9 bolster the superiority of RF over SVM and ANN in timely classification of electric grid faults. Notably, RF required under 2 seconds of additional data beyond the first outlier to identify the fault type of a time series with maximum accuracy. This is a significant enabler for real-time fault prediction.

## 6.4. Intra-class clustering

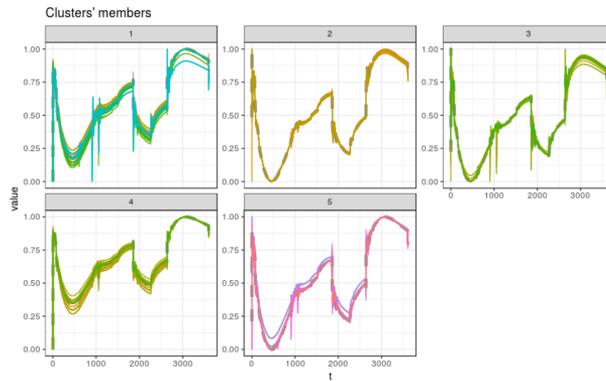
In this stage, we clustered the data within each fault class to uncover additional structure and insight using DTW as the distance metric and PAM as the clustering technique. We show results for the GMD 2 fault class. The inputs to the clustering

were min-max normalized time series. We computed the average intra-cluster DTW distance for different numbers of clusters, plotted in Figure 10. We note from this plot that intra-cluster distance initially drops as the number of clusters increases, and then flattens beyond the case of 5 clusters. We thus determined that the optimal number of clusters to capture structure in the data is 5. The five clusters contained 28, 10, 17, 17, and 54 temporal waveforms, with respective average intra-cluster DTW distances of 17.870670, 5.740929, 14.639842, 17.140088, and 7.272719.



**Figure 10.** Average intra-cluster DTW distance (averaged across all clusters) as a function of the number of clusters.

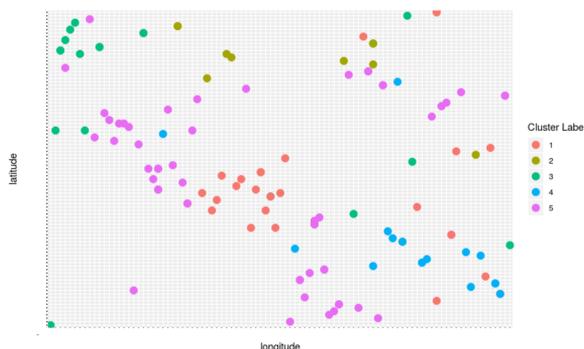
The five clusters contained 28, 10, 17, 17, and 54 time series, with an average intra-cluster DTW distance of 17.870670, 5.740929, 14.639842, 17.140088, and 7.272719 respectively. Figure 11 shows the normalized time series in each cluster.



**Figure 11.** Clustering of GMD 2 fault time series (time series are min-max normalized).

The time series appear to be sufficiently different across clusters. Although there is some intra-cluster variation, splitting the data into additional clusters may cause overfitting.

Figure 12 plots the geographical coordinates of the GMD 2 faults, color-coded by cluster membership. Interestingly, we note that the clusters appear to correspond closely to the geographical distribution of the grid stations, suggesting



**Figure 12.** Geographical representation of clustering of GMD 2 fault time series.

that location affects the response of a station to a fault. Further investigation is required to uncover the deeper significance of these intra-class subgroups in effectively responding to faults, as well as any potential connection to fault severity.

## 7. CONCLUSIONS AND FUTURE WORK

An early warning system prototype to detect, classify, and diagnose BPA electric grid faults was introduced. An efficient and effective anomaly detection method using quartiles was employed to detect the presence of a fault. Following this, the autocorrelation function was used to generate a compact, highly discriminative feature representation of a time series. Comparing the support vector machine, random forest, and artificial neural network classifier models, the random forest classifier was observed to be the superior technique in terms of accuracy, consistency, robustness to sparse training data, and excellent performance on incomplete time series for preemptive fault classification. Finally, intra-class clustering was conducted using partitioning around medoids with the distance metric defined by dynamic time warping. This uncovered five clusters that correlated closely with geographic location.

Future work will focus on gaining deeper understanding of the structure and implications of intra-class subgroups, including the possibility of fault severity estimation. Variables thus far unused, namely voltage and phase angle can be used to provide additional information and granularity in the clustering process. In addition, developing more effective spatiotemporal visualizations of time series in the BPA grid is of crucial importance in enabling operators to assess and act quickly in potentially disastrous situations. These may include real-time heat maps of outlier frequency, rate of change, and other metrics. Additional data such as weather related parameters (temperature, humidity, etc.) could also be incorporated to aid the classification of faults with near-identical time series.

**Acknowledgments.** The authors would like to thank Joyce Luo, Jeremy Nguyen, and Zachary Yung for their contributions to early stages of this work. Thanks to Ms. Erica Bailey at Pittsford-Mendon High School for initiating this

collaboration. Finally, thanks to Professor Esa Rantanen from the RIT department of Psychology for introducing us to the BPA data and the themes therein.

#### REFERENCES

- [1] M. Farshad and J. Sadeh, *Fault locating in high voltage transmission lines based on harmonic components of one-end voltage using random forests*, Iranian Journal of Electrical & Electronic Engineering **9** (2013), 158–166.
- [2] A. N. Hasan, P. S. P. Eboule and B. Twala, *The use of machine learning techniques to classify power transmission line fault types and locations*, 2017 International Conference on Optimization of Electrical and Electronic Equipment (OPTIM), IEEE, 2017, pp. 221–226.
- [3] R. Malhotra and A. Jain, *Fault prediction using statistical and machine learning methods for improving software quality*, Journal of Information Processing Systems **8** (2012), 241–262.
- [4] E. B. M. Tayeb, *Faults detection in power systems using artificial neural network*, American Journal of Engineering Research **2** (2013), 69–85.
- [5] E. F. Ferreira and J. D. Barros, *Faults monitoring system in the electric power grid of medium voltage*, The 8th International Conference on Sustainable Energy Information Technology (SEIT 2018), Procedia Computer Science **130** (2018), 696–703.
- [6] T. Dalstein and B. Kulicke, *Neural network approach to fault classification for high-speed protective relaying*, IEEE Transactions on Power Delivery **10** (1995), 1002–1011.
- [7] M. Jamil, S. K. Sharma and R. Singh, *Fault detection and classification in electrical power transmission system using artificial neural network*, SpringerPlus **4**:334 (2015), 13 pp.
- [8] B. Bhattacharya and A. Sinha, *Intelligent fault analysis in electrical power grids*, CoRR, abs/1711.03026, 2017.
- [9] E. M. Rantanen, E. Fokoué, K. Gegner, J. Haut and T. J. Overbye, *Data properties underlying human monitoring performance*, Proceedings of the Human Factors and Ergonomics Society Annual Meeting **61** (2017), 1711–1715.
- [10] V. N. Vapnik, S. E. Golowich and A. J. Smola, *Support vector method for function approximation, regression estimation and signal processing*, in: M. I. Jordan, M. C. Mozer and T. Petsche (eds.), Advances in Neural Information Processing Systems **9**, MIT Press, 1997, pp. 281–287.
- [11] L. Breiman, *Random forests*, Machine Learning **45** (2001), 5–32.
- [12] C. M. Bishop, *Neural Networks and Pattern Recognition*, 1st ed., Oxford University Press, 1995.
- [13] H. Sakoe and S. Chiba, *Dynamic programming algorithm optimization for spoken word recognition*, IEEE Transactions on Acoustics, Speech, and Signal Processing **26** (1978), 43–49.
- [14] L. Kaufman and P. J. Rousseeuw, *Clustering by Means of Medoids*, Reports of the Faculty of Mathematics and Informatics, 1987.
- [15] S. Lloyd, *Least squares quantization in PCM*, IEEE Transactions on Information Theory **28** (1982), 129–137.
- [16] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2019, <https://www.R-project.org>.
- [17] S. Dégerine and S. Lambert-Lacroix, *Characterization of the partial autocorrelation function of nonstationary time series*, Journal of Multivariate Analysis **87** (2003), 46–59.
- [18] P. Stoica and R. Moses, *Spectral Analysis of Signals*, Prentice Hall, 2005.

Sanjeev Raja, College of Engineering, University of Michigan, Ann Arbor, MI 42111, USA  
*e-mail*: [sanjeevr@umich.edu](mailto:sanjeevr@umich.edu)

Ernest Fokoué, School of Mathematical Sciences, Rochester Institute of Technology, 85 Lomb Memorial Drive, Rochester, New York 14623, USA  
*e-mail*: [epfeqa@rit.edu](mailto:epfeqa@rit.edu)