PNNL-33548

Big Data Analysis of Synchrophasor Data

Outcomes of Research Activities Supported by DOE FOA 1861

October 13, 2022

Shuchismita Biswas
Jim Follum
Eric Andersen
Jeffery S Banning

Prepared for the U.S. Department of Energy
Under contract DE-AC05-76RL01830

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <https://www.ntis.gov/about>
Online ordering: http://www.ntis.gov

PNNL-33548

Big Data Analysis of Synchrophasor Data

Outcomes of Research Activities
Supported by DOE FOA 1861

October 13, 2022

Shuchismita Biswas
Jim Follum
Eric Andersen
Jeffery S Banning

Prepared for
the U.S. Department of Energy
Under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Abstract

This report describes the key outcomes of research activities sponsored by the Department of Energy's Funding Opportunity Announcement (FOA) number 1861 that was aimed at advancing the state-of-the-art in big data analytics applied to transmission-level synchrophasor measurements. The FOA resulted in eight research grants where the awardees developed machine learning and artificial intelligence tools and approaches. The commonalities in tools and approaches used by the awardees are explored, and insights gained from how the project outcomes might be operationalized are discussed. This report does not seek to comprehensively summarize all research supported by the FOA, rather it focuses on enabling the fast dissemination of major findings to the broader power systems community.

Acknowledgments

The authors would like to thank Sandra Jenkins of DOE's Office of Electricity and Carol Painter of the National Energy Technology Laboratory (NETL) for their administration of the FOA 1861 effort. The authors are also grateful to the organizations that contributed synchrophasor measurements to build the FOA 1861 dataset.

Acronyms and Abbreviations

AE	Auto-Encoder
AEM	Automated Event Miner
AI	Artificial Intelligence
ANN	Artificial Neural Network
AWS	Amazon Web Services
BDA	Big Data Analytics
BTrDB	Berkeley Tree Database
CNN	Convolutional Neural Network
CPU	Central Processing Unit
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
DOE	Department of Energy
DTW	Dynamic Time Warping
EPG	Electric Power Group
FOA	Funding Opportunity Announcement
GAN	Generative Adversarial Network
GPS	Global Positioning System
GPU	Graphics Processing Unit
GSP	Graph Signal Processing
HDFS	Hadoop Distributed File System
IoT	Internet of Things
LSTM	Long Short-Term Memory Networks
LTTB	Largest Triangle Three Buckets
ML	Machine Learning
MTF	Markov Transition Field
PCA	Principal Component Analysis
PMU	Phasor Measurement Unit

PNNL	Pacific Northwest National Laboratory
RAS	Remedial Action Scheme
SEL	Schweitzer Engineering Laboratories
SIN	Sparsity Inducing Norm
SNR	Signal-to-Noise Ratio
SPP	Spatial Pyramid Pooling
SQI	Stream Quality Index
SVM	Support Vector Machine
TRL	Technology Readiness Level
UTC	Coordinated Universal Time
WAMS	Wide Area Measurement System

Contents

Figures

Tables

1.0 Introduction

With the aim of advancing big data-driven research in power systems, the Department of Energy (DOE) created Funding Opportunity Announcement (FOA) 1861, wherein a massive dataset of synchrophasor measurements and event logs spanning multiple years and the three electrical interconnections in the continental United States was collated, anonymized, and distributed to eight FOA awardees [1–3]. An assessment of how effective the data anonymization process was and how it impeded research is provided in [4]. The awardees were not aware of the identities of the data contributors, the location of Phasor Measurement Units (PMU), or the underlying transmission network topology. Nevertheless, their research efforts highlighted how actionable insights can be derived from the massive archives of PMU data that utilities have, and how Big Data Analysis (BDA), Artificial Intelligence (AI) and Machine Learning (ML) methods can facilitate this process. This report highlights major outcomes and insights from the research activities supported by DOE's FOA 1861, identifies commonalities in the approaches adopted by the different project teams, and discusses how to operationalize the lessons learned. Key outcomes include-

- Robust data management platforms were built for PMU big data analytics at scale. The project teams devised effective methods for cleaning, visualizing and performing computations on over 20 TB of PMU data. Techniques for 'streaming' data from the archive were also developed to emulate how the proposed algorithms would perform in near real-time environments.

- The performance of different event detection and classification algorithms at varying Technology Readiness Levels (TRL) were demonstrated- ranging from novel academic approaches to commercially available proprietary tools. A byproduct was the formulation of several feature engineering methods- from statistical description to image encoding techniques, that may prove useful for AI/ML applications beyond event detection and classification.

- Besides the events present in the utility logs, the awardees also detected and classified numerous additional events, often with the help of subject matter experts. Open-source synthetic event data was also created and made publicly available. As these events come from geographically diverse systems, signatures derived from this data repository will be broadly applicable to different power systems, and hence will help evaluate and improve the performance of existing commercial situational awareness tools. The event signatures can further be incorporated into educational tools for training students and operators.

- Transfer learning techniques were developed to extend models trained with data from one interconnection to another. Methods to augment measurement data with simulations for the purpose of training ML models were also shown for events whose field occurrences are rare. These developments are notable because they show that models trained on a system with good data quality can be deployed on systems with sparse historical records with relative ease, and without having to start from scratch.

- The awardees also pursued several other research directions, such as looking for evidence of precursors to disturbances and equipment failures, and mitigating GPS signal spoofing. Although data anonymization and the lack of detailed labels and system knowledge made validating these results challenging, the developed approaches have merit. System operators can choose to evaluate these approaches with data from their footprints and pursue further refinement if they prove promising.

This report does not seek to comprehensively summarize all research supported by the FOA, but is rather aimed at facilitating the fast dissemination of the major findings to the broader power systems community. Moreover, the authors have not independently evaluated the performance/accuracy of the algorithms beyond what has been self-reported by the awardees. Again, the intention of this report is not to judge which algorithm is the ’best’, but rather concisely lay out the developed approaches and their suitability for various applications.

For convenience, the following nomenclature will be used to refer to the awardee projects in the remaining report.

Table 1. Awardee project nomenclature

Reference	Lead	Team Members	Project Title
Project 1	Ping Things		Combinatorial Evaluation of Physical Feature Engineering and Deep Temporal Modeling
Project 2	GE Research	GE Grid Solutions	PMU-Based Data Analytics using Digital Twin and PhasorAnalytics Software
Project 3	Schweitzer Engineering Laboratories	Oregon State University	Machine Learning Guided Operational Intelligence from Synchrophasors
Project 4	Siemens Corporation	Southern Methodist University, Temple University	MindSynchro
Project 5	University of California, Riverside	Electric Power Group (EPG), Michigan Technological University	Discovery of Signatures, Anomalies, and Precursors in Synchrophasor Data with Matrix Profile and Deep Recurrent Neural Networks
Project 6	University of Nevada, Reno	Arizona State University, IBM, Virginia Tech	A Robust Event Diagnostics Platform: Integrating Tensor Analytics and Machine Learning Into Real-time Grid Monitoring
Project 7	Iowa State University of Science and Technology	Electric Power Group (EPG), Google Brain, IBM	Robust Learning of Dynamic Interactions for Enhancing Power System Resilience
Project 8	Texas A&M Engineering Experiment Station	Temple University, Quanta Technology	Big Data Synchrophasor Monitoring and Analytics for Resiliency Tracking (BDSMART)

A list of technical publications from each project is provided in Appendix A. Readers are encouraged to refer these publications for further details of methodologies only briefly described in this report.

The rest of the report is organized as follows. Chapter 2.0 describes data management platforms used by the awardees, and Chapter 3.0 summarizes some of the algorithms developed. Some extracted event signatures are presented in Chapter 4.0. Chapter 5.0 presents some initial thoughts on how to operationalize key findings from the research enabled by FOA 1861 and concludes this report.

2.0 Data Management

The PMU dataset was delivered to the awardees as snappy-compressed partitioned parquet files on SATA hard drives. Data was provided in two installments- about 18 TB of *training* data was initially delivered to train models, followed by about 7 TB of *test* data to evaluate their performance on unseen measurements. The awardees had to develop efficient computation pipelines to process this massive amount of data and discover patterns present within. This chapter provides a high-level overview of the various tools used to realize these pipelines.

2.1 Data Storage

Two main approaches were used for storing the raw data- a) cluster-based technologies, and b) cloud-based services, as tabulated below.

Table 2. Data Storage Methods

Approach	Tools Used
Cluster-based	Data is stored on servers using Apache Hadoop framework (Project 2, 5)
Cloud-based	Amazon S3 (Amazon Simple Storage Service) (Project 3) IBM COS (Cloud Object Storage) (Project 7)

Hadoop based frameworks for storing and processing big data are popular due to their high throughput and fault tolerant processing performance. Data is stored using the Hadoop Distributed File System (HDFS) and queried using Apache Hive. Cloud-based storage services are becoming increasingly popular particularly as they also enable easily scalable high-performance downstream compute environments.

Project 1 used their proprietary PingThings PredictiveGrid platform to ingest, store and analyze data in the Berkley Tree Database (BTrDB) that the platform is built on. In the tree structure of the database, internal nodes store summary statistic for varying durations of the raw data stored at the leaf nodes. A visual representation of the BTrDB structure is shown in Fig. 1.

2.2 Speeding Up Computation

Most projects used Python and its popular packages to develop their ML pipelines and leveraged various parallelization techniques to speed up computations. The algorithms in Project 1 were built on the proprietary PredictiveGrid Platform, and benefited from the tree-like structure of their database that automatically computed and stored summary statistics at different temporal resolutions. Several other projects like Project 2 and 3 also stored minute-level summary statistics of the raw data during the initial data preprocessing stages to accelerate downstream computations.

Project 2 developed a scalable feature generation pipeline using Spark v3.0.0 where all computations were carried out in a distributed manner over cluster servers. Spark's mapInPandas transformation was used to express feature generation algorithms as map functions that get invoked by each executor. This avoided the need to convert Spark dataframes to Pandas dataframes (a notoriously slow process due to the data serialization/deserialization overheads) and then carrying out feature generation computations sequentially on a single server. Hence, computation time was greatly reduced.

GPU-based accelerations were used by several projects. For instance, Project 3 concluded that GPU-based accelerations were particularly advantageous over distributed computing on

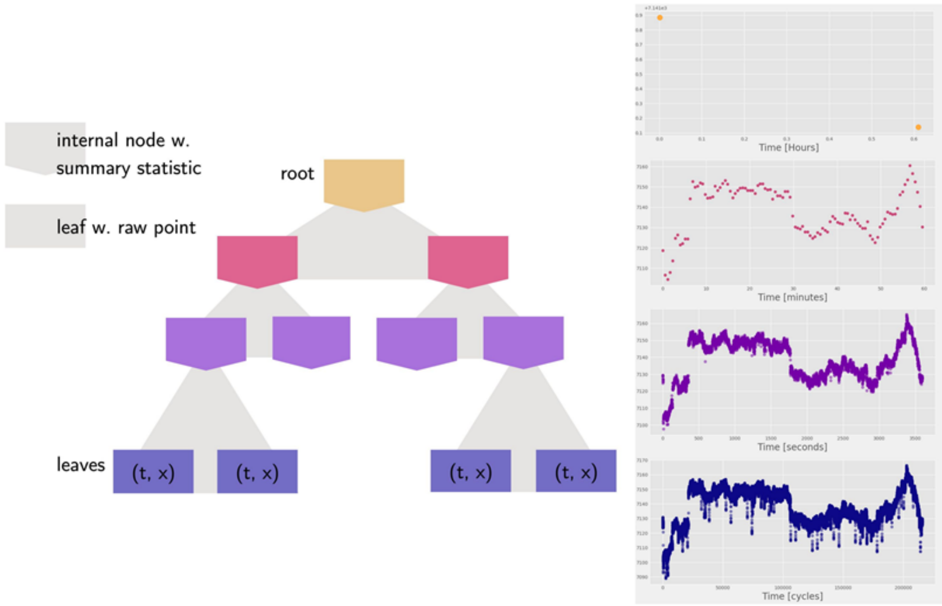

Figure 1. The structure of BTrDB, the database used in Project 1

CPUs as the number of features considered and algorithm complexity increased. Project 4 utilized RAPIDS cuDF, a recently released python library that provides pandas-like functionality for performing GPU-based data processing operations, to speed up data preprocessing tasks. Cloud-based algorithm pipelines were also adopted. Some examples are:

- Project 3 implemented a distributed computation framework using Amazon Web Services (AWS) Elastic Compute Cloud (EC2) instances controlled by Ansible to run their disturbance detector over the training dataset. Multiple EC2 instances were created; each instance fetched a predetermined data chunk from Amazon S3 storage and ran the detector over it. Metadata from the detected events were saved and stored back in Amazon S3.

- Project 7 used Kubernetes, an open-source system for automating the deployment, scaling, and management of containerized applications. The project team containerized the developed ML libraries using Docker, hosted data on IBM COS, and utilized the Ray distributed computing framework and OpenShift clusters to deploy their ML pipeline. They noted a tenfold improvement in computation time with this cloud-based framework over their previous cluster-based solutions.

- An application was created on the MindSphere cloud platform to perform inferences based on semi-supervised ML models already trained in Project 4. MindSphere is the industrial IoT (Internet of Things) as a service offering from Siemens, and can be employed in industrial applications to ingest and visualize real-time data and analytics results in one centralized location with minimal development effort.

The project teams explored various resources before zeroing in on their final data handling pipelines. Their experiences and lessons learned (particularly which tools did not work well or were not mature enough when FOA 1861 activities were conducted) are shared in greater detail in their final reports. This will prove to be a great resource for organizations looking at building their own BDA tools and pipelines for analyzing PMU data.

3.0 Developed Algorithms

All awardee projects developed and tested event detection and classification algorithms that directly supported the stated FOA objectives of discovering key events and unlabeled disturbances present in the dataset for the purpose of deriving characteristic event signatures. Some of the proposed algorithms warrant parallel implementation on individual PMUs, while others are designed to utilize measurements from multiple PMUs simultaneously. Moreover, to support the event detection and classification pipelines, several data preprocessing and feature-engineering approaches were developed and may find application beyond the use-cases demonstrated in the current research efforts.

Data Cleaning (Section 3.1)	Bad Data Detection	- Metrics: Stream Quality Index (SQI) (Project 1) - Rule-based detection of outliers, stale data and physically impossible values - Using PMU status bit information
	Imputation	- Simple linear regression (Project 3) - Event-participation decomposition model SPIKE-P (Project 5) - Attentional LSTM models (Project 5) - Regularized low-rank tensor completion (Project 6)
Feature Engineering (Section 3.2)	Statistical measures	- Sample mean, minimum, maximum and standard deviation - Signal-to-noise ratio, Kurtosis, skewness - Quantiles, range, coefficient of variation - Auto-regressive measures - Ramp-up and ramp-down rates (Project 6) - Amplitude above and below sample mean (Project 6) - Area above and below sample mean (Project 6)
	Signal-processing based	- Continuous Wavelet Transform (Project 1) - Linear Predictive Coding (Project 1) - Prony analysis (Project 3,4) - Spectral quantities (Project 4)
	Miscellaneous	- Critical slowing down (Project 3) - Markov transition field (MTF)-based image encoding of time-series (Project 7) - Soft dynamic time warping (Project 8)
Event Detection (Section 3.3)	Statistical/ Mathematical	- Unsupervised anomaly detection using statistical measures (Project 1,2) - Prony analysis-based ringdown detector (Project 3) - Rule-based method checking rank signatures of PMU data matrices (Project 6)
	ML-based	- Semi-supervised Principal Component Analysis (PCA)-based normality modeling (Project 2) - Critical slowing down+ANN (Project 3) - Supervised graph signal processing-based method (Project 5) - Supervised sparsity-inducing norm-based method (Project 5) - Supervised bidirectional anomaly Generative Adversarial Network (BiAnoGAN) (Project 5) - Supervised CNN-based models with automated denoising (Project 8)
	Commercial tools	- SEL SynchroWave (Project 3) - EPG Automated Event Miner (AEM) (Project 5,7)
Event Classification (Section 3.4)	Unsupervised	- Clustering followed by SME labeling (Project 4)
	Supervised	- Random Forests (Project 1,3,6) - CNN using GSP-based sorting and information loading-based regularization (Project 5) - Spatial pyramid pooling (SPP)-based CNN (Project 7) - Deep graph learning using interaction-graphs (Project 7) - Hierarchical CNN-based models (Project 8)
	Semi-supervised	- AE-based models (Project 4) - CNN-based safe tri-net method (Project 7)
Data Visualization (Section 3.5.2)		- Mr. Plotter (Multiresolution plotter) (Project 1) - Largest Triangle Three Buckets (LTTB) (Project 8)

Figure 2. Thematic grouping of some algorithms used by the projects supported by FOA 1861

Some of these methods are thematically grouped and summarized in Fig. 2, which highlights the commonalities in approaches adopted by the various awardees. The rest of this chapter presents further details about these methodology groups. Techniques commonly used by almost all projects have not been attributed to individual awardees in this summary.

3.1 Data Cleaning

The varying data quality across PMUs in the dataset posed major challenges to algorithm development. Besides obvious sample drops, the data exhibited myriad issues like quantization noise, inconsistent usage of units, stale values, and timestamp jitter. Nevertheless, the project teams devised methods to detect and fix data quality issues before feeding the data to their downstream analysis pipelines.

3.1.1 Bad Data Detection

All the teams used some rule-based methods for detecting statistical outliers, physically impossible measurements (negative voltage magnitudes, voltage/current angles outside $\pm 180°$ etc) and stale values (measurements repeated for longer than a specified threshold), and removing them from the data. As many of these rules could be applied using only summary statistics, identifying regions of bad data was computationally quite efficient. Project 1 formulated a composite metric called the stream quality index (SQI) to formally express the 'goodness' of a PMU data stream. This SQI metric looked at five individual metrics- point count density, deviation of measurements from base kV values, obviously bad values, repeated values and timestamp jitter.

Status bits reported by PMUs can also be used to identify samples where time synchronization and other data quality issues are present. Several projects effectively used the PMU status bits to cleanse bad data samples. However, it must be noted that measurements may be erroneous even when not flagged by the status bits.

3.1.2 Data Imputation

As many ML algorithms cannot handle missing data points, data imputation methods become necessary. In many cases, when only a few intermittent samples are missing, linear interpolation may be sufficient. In other cases with large swathes of consecutively missing data, devising more sophisticated methods becomes necessary. Project 3 found that the simple linear regression model was able to reconstruct missing data for one PMU using information from other PMUs in the system, and the performance was best when using information from highly correlated neighbors. The technique struggled the most with matching the DC trend of the data, which can be expected as the state of the power system is always changing.

Project 6 utilized the low-rank nature of PMU data to devise a regularized tensor-completion based data imputation method. By stacking all measurement channels of multiple PMUs into tensors, this method is able to utilize the spatial and temporal correlations present across PMU channels for effective data reconstruction. Data inter-dependencies and system constraints were used as regularization terms in the low-rank tensor completion problem to restrict the search-space to meaningful solutions and improve reconstruction accuracy.

Project 5 devised ML-based methods for data recovery. The SPIKE-P (Stochastic and Proximal Implicit Krasulina Event-Participation) algorithm is designed for online imputation of streaming event data- events are decomposed into non-dynamic components that represent

PMU participation in an event (derived from past data), and dynamic components that represent disturbance magnitude (derived from current timestamp). The algorithm shows high accuracy and fast performance. Their attentional LSTM model is also similarly aimed at online forecast of incoming measurements.

3.2 Feature Engineering

The awardees found that feeding features engineered from the raw data to ML models had advantages over feeding in the raw data itself, both in terms of accuracy and computational performance. A variety of feature-engineering methods have been used in the projects, including simple statistical descriptors over data windows, measures derived from the frequency-domain, and more sophisticated image-encoding techniques. As these features may be correlated, some projects leveraged dimensionality reduction techniques like Principal Component Analysis (PCA) to reduce the number of features to be fed to ML models.

Most projects utilized statistical descriptors like sample mean, median, standard deviation, etc., to some extent in their feature-engineering pipelines. Features based on Prony analysis (a method of analyzing ringdown oscillations) were constructed using proprietary tools and utilized in Projects 3 and 4 with varying degrees of success. The statistical and signal-processing based features have the advantage of easy interpretability- it is intuitive to attribute certain statistical/spectral behavior to particular disturbances. On the other hand, image-encoding of time-series data, while less interpretable, can help identify complex interactions present in the data that is not otherwise easily discernible. Dynamic time warping (DTW) can help ML models recognize similarities in temporal sequences that are not perfectly aligned.

A detailed list of the features is available in the project reports prepared by the awardees (Appendix A). Although the features were devised for particular ML pipelines, they can be adopted for training models aimed at other applications.

3.3 Event Detection

Methods with varying TRLs were used for detecting events- from experimental deep neural networks to commercially available tools.

3.3.1 Statistical and Mathematical Approaches

Some projects used summary statistics to identify anomalous periods in the data, followed by voting strategies that accounted for how many metrics detected the anomaly and how many PMUs 'saw' it in order to determine whether the anomaly was indeed an event. As digging through the summary statistics was computationally much easier than running computations with the raw data, awardees utilized these methods as workhorses to mine disturbances present in the entire dataset. More computationally intensive methods were illustrated using smaller chunks of the raw data.

Project 6 used the low-rank property of PMU data in normal operating conditions to devise a detection algorithm. The idea is that the low-rank property does not hold during disturbances, and hence if the rank crosses a certain threshold, events may be present in the system. A bayesian optimization based parameter tuning method was adopted to optimally choose threshold values. The performance of this event detection method may suffer when deployed in smaller systems where PMUs may exhibit highly correlated response to disturbances (like generator loss outside area), thereby not increasing the rank value significantly.

3.3.2 Commercial Tools

Commercial tools utilized included Electric Power Group's (EPG) Automated Event Miner (AEM) and SEL's SynchroWave Event Detector. Both of these tools were very effective in identifying thousands of disturbances from the two years of data, many of which had not been reported in the utility-provided event logs. While the AEM tool was originally designed to mine historical records, the SynchroWave platform is aimed at processing streaming data. Hence, the SynchroWave tool had to be refactored for processing the FOA 1861 dataset.

3.3.3 ML models

Different supervised and semi-supervised ML models were trained for detecting power system events. Using semi-supervised models was necessary due to inconsistent and incomplete labeling in the utility-provided event logs.

In Project 2, a PCA decomposition-based method was used to fit a normality model. Normality models are semi-supervised ML models trained from data in only one class: the normal class. Once trained, they can predict or assess anomalies within new data, and also estimate feature contributions to anomalies. Instances of normal operations were chosen from periods outside the vicinity of known events present in the event logs. Each data example fed to the normality model was the computed feature vector for one PMU at one timestamp. Based on the trained model, anomaly scores for new data samples were calculated. Estimating which features contributed to particular types of anomalies was beneficial for streamlining the performance of their event classification models. The project had also investigated the use of autoencoders, but found the PCA method to be better-suited for the application.

Project 5 explored several supervised deep learning models for the event detection task. It is envisioned that the models will be trained offline, and then deployed online for detecting occurrences in streaming PMU data. A graph signal processing (GSP)-based algorithm was developed to effectively utilize spatial and temporal correlations present in streaming PMU data for anomaly detection. The algorithm has linear time complexity in the online environment, making it extremely scalable. The sparsity-inducing norms (SIN)-based method was designed specifically for detecting voltage events. Here, the intuition that the low-rank property of PMU data does not hold true during disturbances was utilized (the same principle utilized in the rank-signature based detection method proposed by Project 6). PMU data matrices were decomposed into a low-rank matrix, a sparse event pattern matrix, and a noise matrix. Features computed from these matrices were then clustered using the DBSCAN (Density-Based Spatial Clustering of Applications with Noise) algorithm to identify the presence of anomalies. Project 5 also developed a generative adversarial network (GAN) based event detection model called BiAnoGAN that could be trained using only a few instances of labeled events. The BiAnoGAN was more selective than the GSP-based method, and could detect more than just voltage events unlike the SIN-based method.

Project 8 developed several convolutional neural network (CNN)-based event detection pipelines utilizing automated feature-learning and data denoising. The CNN-based methods were found to outperform traditional benchmark algorithms such as logistic regression, and support vector machines (SVM). Several awardees reported that it was easier to build standard learning-based models (decision trees, SVM, CNN, etc.) for PMU data-driven applications. Some projects had difficulty tuning the parameters for state-of-the art deep learning models like Auto-Encoders (AE) and GANs, and hence could not achieve optimal performance.

3.4 Event Classification

Binary vs. Multi-class Classification

As power system events may be categorized into multiple types and sub-types, an important question arises here- should one use a multi-class classifier or multiple binary classifiers? Awardees used both approaches - some developed multi-class classifiers (Project 1,5,6,8) while others trained an ensemble of binary classification models (Project 2,3). Ensemble learning with multiple binary event classifiers was observed to work well, whereas the performance of multi-class models was degraded when the number of target classes increased too much (Project 1). Good performance was obtained from 3-class and 5-class models (Project 5,6,8).

Hierarchical Classification

Project 4 and Project 8 explored the use of hierarchical classification in place of conventional multi-class variants. In Project 4, the concept of hierarchy was used to capture the interdependency between event labels. For instance, line trips with and without short circuits are both line trip events and may be categorized as sub-classes within the 'line trip' class. In the approach taken by Project 8, data was first determined to be either normal or anomalous, and anomalous data windows were further classified into line faults and frequency events. Such hierarchical classification presents a promising opportunity for adding granularity to power system event classification models.

Use of CNN Architecture

The CNN architecture was a popular choice among the awardees, with most projects utilizing it in some form in their deep learning pipelines. This is because in recent years, CNN-based models have shown superior performance in image classification and object identification tasks in the computer vision domain. Multichannel PMU data can be stacked into image-like 2D matrices or 3D tensors making them suitable to be fed into CNN-based pipelines. Project 5 utilized such a CNN-architecture and sought to improve its performance by using GSP-based sorting of highly correlated PMUs and information-loading based regularization. Project 7 used the Markov Transition Field (MTF) technique to encode temporal PMU data into 2D images to be used to train a CNN-based classifier. In this pipeline, Spatial Pyramid Pooling (SPP) was used to account for mismatches in data lengths.

3.5 Supporting Tools

In addition to end-to-end event classification pipelines, the awardees also formulated other tools to facilitate data exploration and improve the performance of their proposed algorithms. Some of these tools are described next.

3.5.1 Automated Labeling Applications

To obtain reasonable performance from their detection and classification models, most awardees felt the need to refine the utility-provided event logs. Many of them relied on domain experts to either manually label events or validate labels assigned by automated applications. A clustering-based approach was devised in Project 4- all detected events were divided into several clusters using the k-means algorithm, and then domain experts inspected patterns from each cluster and assigned event labels if the patterns were observed to be distinct. In Project 2, the anomaly scores assigned by the normality model to event data were used to refine the temporal precision of event logs. A heuristic temporal localization algorithm was also developed in Project 1. The automated label improvement strategies thus developed can also help other utilities in extracting event records from their PMU archives. These methods can significantly reduce engineering staff burden, as domain expertise is a prerequisite for reliably creating event labels. Unlike other deep learning applications such as image classification, the labeling process cannot be crowd-sourced.

3.5.2 Data Visualization Techniques

Data visualization plays a critical role in data-driven knowledge discovery, especially at the data exploration stage. Some tools used by the awardees for data visualization include-

- *Mr. Plotter, or Multiresolution Plotter* (Project 1): This tool, built on the PredictiveGrid platform, retrieves and plots statistical aggregates of data (mean, min, max, count) stored in the BTrDB database, which alleviates bottlenecks in fetching high volumes of raw data when rendering plots. The result is a visualization tool that allows users to interactively visualize time-series data containing up to millions of data points using an application that automatically queries new data and intelligently adjusts the plot as users zoom and scroll. The tool allows for seamless visual exploration of data across different time scales, thereby enabling easy identification of periods of interest for further analysis.

- *Largest Triangle Three Buckets (LTTB)* (Project 8): LTTB is a downsampling technique, aimed at preserving visual similarity to the raw data. Instead of removing uniformly spaced data points, LTTB only removes visually redundant data points. This results in small file sizes, thereby enabling fast querying and plot rendering. The potential for using LTTB in feature engineering for ML pipelines has also been explored in Project 8.

3.5.3 Transfer Learning

Transfer learning refers to repurposing pre-trained models for new tasks or applications. Projects 5 and 8 demonstrated that using transfer learning techniques, event detection and classification models trained for one interconnection can be deployed to other interconnections without significant drop in model accuracy. This highlights that ML algorithms trained on systems with good data availability can be generalized to other power systems that may have limited or inconsistent event labels. Project 8 also explored how simulation data can be used to augment field measurements for training ML models, while mitigating issues like class imbalance.

3.5.4 Synthetic Data Generation

In order to address the lack of labeled data availability impeding the advancement of data-driven power systems research, Project 5 used data-driven generative models to generate almost 1000 instances of labeled synthetic data. This synthetic dataset (named pmuBAGE) has been made publicly available to encourage benchmarking the performance of developed ML algorithms against a standard dataset. This will prove to be a valuable resource for researchers and developers who do not have access to field measurements.

3.6 Approaches Requiring Further Research and Validation

In addition to the algorithms described so far, the FOA awardees also pursued applications aimed at identifying complex patterns present in PMU data. Although the methods had merit, some approaches were inconclusive and achieved limited success due to limitations in the dataset itself. For instance, some findings could not be validated in the absence of detailed event labels. As no information about the topology and spatial proximity of the PMUs was available to the awardees, use-cases like fault localization could not be pursued and the full potential of approaches such as graph neural networks could not be realized. Measures used to overcome event log limitations included:

- Many projects used the spatial spread of events to judge which were more severe and hence of greater interest to the utilities. However, in the absence of topological information, correctly inferring this metric was a challenge. Some awardees like Project 1 used the number of PMUs that detected an event as a proxy for the event's spatial spread. This approach could be misleading as PMUs in the dataset were unevenly distributed- some regions had dense PMU clusters, whereas other areas had a relatively sparse distribution. Other awardees used correlation studies to infer proximity information about the PMUs- Project 2 used correlations in signal-to-noise ratio (SNR) in voltage magnitude channels, while Project 3 utilized voltage dip proportions during faults.

- Detecting early warning signs of equipment failure is a promising application of PMU data. However, examples of properly labeled equipment failures were rare in the event log. Awardees used their interpretations of the log contents to identify possible failure cases, but their findings could not be validated in the absence of additional information.

- Event start times noted in the event logs were imprecise because for most utility uses only the general timing of the disturbance is needed. Most projects found that their detectors/classifiers did not perform well when the utility-provided timestamps were used directly. Hence, the awardees spent some time refining the event start times to improve the performance of their models. For example, Project 2 utilized their trained normality model to further refine event logs, while Project 1 developed a heuristic localizer called H-loc.

Some other approaches employed in the projects are described next.

3.6.1 Oscillation Analysis

Ringdown events detected in the dataset were analyzed in various ways by the awardees. Project 3 identified the frequencies of the dominant modes activated by the ringdown events and were able to generate unique fingerprints for each of the three interconnections.

Project 2 computed the frequencies and damping ratios of modes activated by the detected ringdown events and investigated whether the modal properties could be correlated with seasonal patterns. Their study was inconclusive. For example, they could not find evidence of negative correlation between damping ratio and system loading conditions in the Eastern Interconnection data. This led them to conclude that the dataset perhaps did not include PMUs which have greater observability of the system modes.

3.6.2 Data Anomaly Mitigation

Project 3 utilized the low-rank property of PMU data to devise a methodology for removing the error introduced by GPS spoofing. A PCA-based method was used to learn the low-dimensional subspace where measurement data was expected to lie. A GPS spoofing attack on a PMU shifts its time reference, thereby introducing a fixed bias to all its phase angle measurements. Such data with anomalies would lie outside the legitimate low-dimensional subspace learned by the PCA-based model, which makes it possible to detect and correct the anomalies by projecting the spoofed data onto the known legitimate subspace. It was demonstrated that the proposed method was able to accurately reconstruct data under simulated fixed and ramped spoofing attacks.

However, GPS spoofing is not the only way phase angle biases may be introduced into PMU measurements. For example, it has been previously reported that leap second adjustments to the UTC introduce such biases, and PMUs by the same manufacturer may face similar issues. The proposed spoofing mitigation algorithm needs to be further evaluated to determine if it will introduce additional artifacts when the time reference is affected by issues like leap second addition and loss of clock synchronization.

3.6.3 Precursor Identification

There is evidence in the literature that the SNR in voltage magnitude measurements exhibit greater variability in the lead up to a transformer failure [5]. Motivated by this, Project 2 sought to find precursors to transformer failures in the FOA 1861 dataset. They used the Granger causality model to identify links between periods of high SNR variability in voltage channels and transformer trips reported in the event log, and found instances where high SNR variability was observed a few hours before a reported trip. It must be noted that the reported trip events might have been caused by factors other than transformer failure, and hence the results from this study might need further validation. The methodology, however, may be a good starting point for investigations by utilities.

Project 1 checked whether one minute of data (precursors) before an event can be used to classify event types. Their experiments showed that their multi-class classifier performance was only slightly degraded when using these precursor windows instead of event data.

4.0 Event Signatures

Following event detection, the awardees invested significant research efforts in going through the event repository and identifying signatures and commonalities in event groups. This understanding of signatures helped them in fine-tuning their event classification strategies. The investigation further yielded certain qualitative observations that not only verified power engineering intuitions, but also proved helpful in illustrating expected system behavior to ML practitioners with little power systems background.

Informed by the utility-provided event logs, the awardees tended to focus on three broad groups of events- frequency, oscillation and voltage, as summarized in Table 3. Qualitative characteristics of the categories are evident-

- Frequency events tend to manifest system-wide, persist for quite long (minute-scale), and are characterized by wide-area sag/swell in frequency and voltage (sag in case of loss of generation, swell for loss of load). These events are followed by step changes in active and reactive power as flows throughout the system are redistributed.
- The timescale of oscillation events vary- ringdowns usually dissipate within tens of seconds while forced oscillations may persist much longer. Ringdowns are seen throughout the system, while forced oscillations may also be localized.
- Voltage events tend to be short-lived, localized and are characterized by spikes in voltages and currents. During short circuit line/transformer faults, there usually is a large spike in current and large dip in voltage caused by a sudden draw of power. If the fault results in a trip, then current on the tripped line goes to zero and network power flows are redistributed.

Table 3. Event categories

Event group	Event type
Frequency	Loss of generation
	Loss of load
Oscillation	Forced oscillation
	Ringdown
Voltage	Line trip
	Transformer trip

The awardees found that utilizing these insights about system behavior changes during events, they were able to select different classification algorithms best-suited for each event group. ML-based models were not particularly successful in detecting/classifying oscillation events, while signal-processing techniques like RMS-energy based detection and Prony analysis fared much better (Project 3,4,5,7). Project 5 developed a sparsity-based detector specifically to look for voltage events, motivated by the fact that their GSP-based detector was good at finding frequency events but not voltage ones.

4.1 Identified Signatures

In this section, signatures for the different event categories described in Table 3 are illustrated with the help of recorded measurements. The signatures were observed to be consistent across interconnections, implying that the developed algorithms can be generalized to other power systems as well. The plots in this section use data from multiple interconnections.

4.1.1 Frequency Events

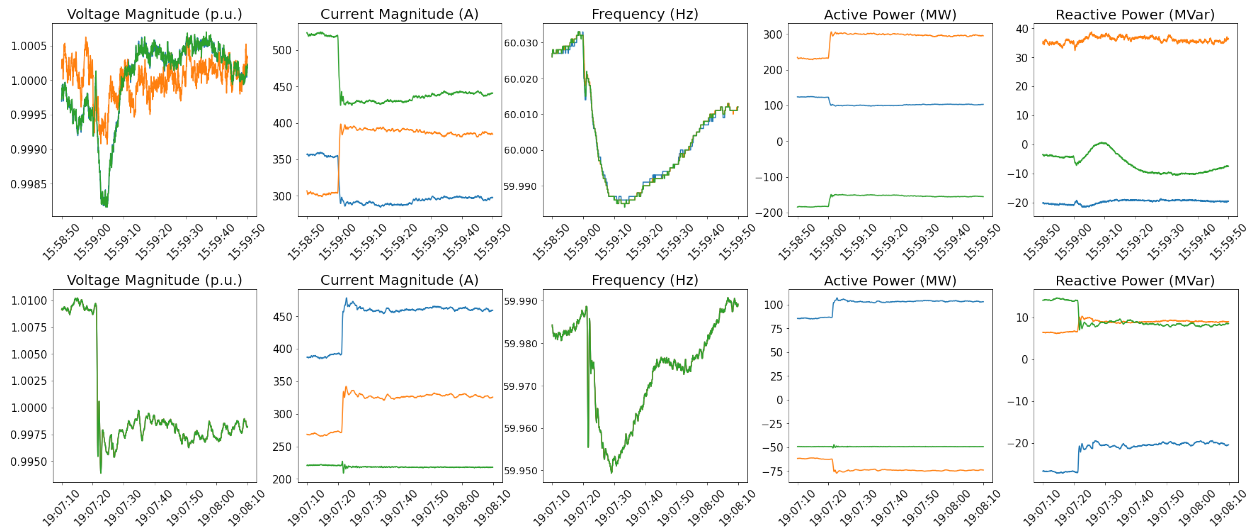

Figure 3. Loss of generation examples. The plots show 1 minute of data.

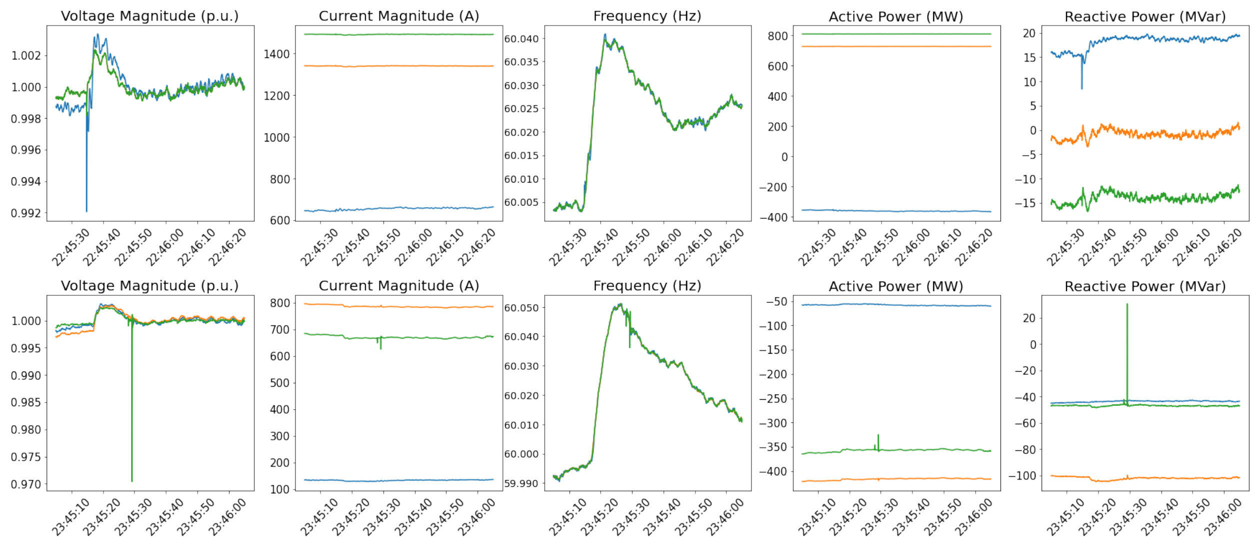

Figure 4. Loss of load examples. The plots show 1 minute of data.

Inspecting the data for frequency events (some examples are shown in Fig. 3 and 4), Project 4 arrived at qualitative signatures for generator and load loss events, listed in Table 4.

4.1.2 Oscillation Events

Forced oscillations are sustained periodic disturbances introduced into the system by some external driving input. They may persist for different time-scales, ranging from seconds to hours. Ringdown oscillations are damped oscillations caused when inter-area modes get

Table 4. Signatures for Generator and Load Loss Events

Event type	Signature
Loss of generation	Voltage drops in some lines and may have oscillations. The voltage of some lines may also increase. Current rises in some lines and may have oscillations. The current of some lines may also decrease. Frequency decreases rapidly and deviates from the nominal frequency. It then recovers to normal after some time ranging from a few seconds to several minutes.
Loss of load	Voltage rises and recovers within a few seconds to several minutes. Current decreases and recovers within a few seconds to several minutes. Frequency increases rapidly due to the imbalance and returns to normal within a few seconds to several minutes. Oscillations may occur.

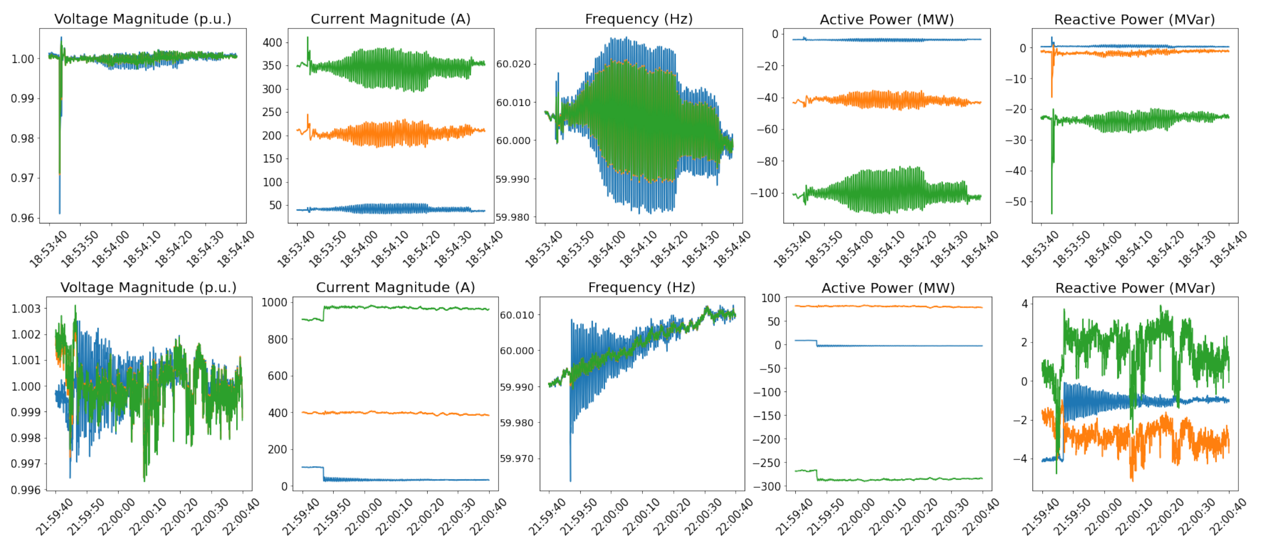

Figure 5. Forced oscillation examples. The plots show 1 minute of data.

excited following a system disturbance such as a generator or line trip. Ringdowns tend to be apparent across wide areas, while forced oscillations can be local or widespread. Some examples of oscillation events from the FOA 1861 dataset are shown in Figs. 5 and 6.

4.1.3 Voltage Events

Voltage events associated with line and transformer trips show similar behavior. As evident from Figs. 7 and 8, the awardees found distinguishing between line and transformer trips to be very challenging. Project 4 summarized the signature for general voltage events as follows:

- Voltage suddenly drops when the short circuit happens.
- Voltage returns to normal within 100ms to 2s.
- Current rises to a very high value and recovers to normal in fractions of a second or drops to zero if it is tripped.
- Power flow changes.

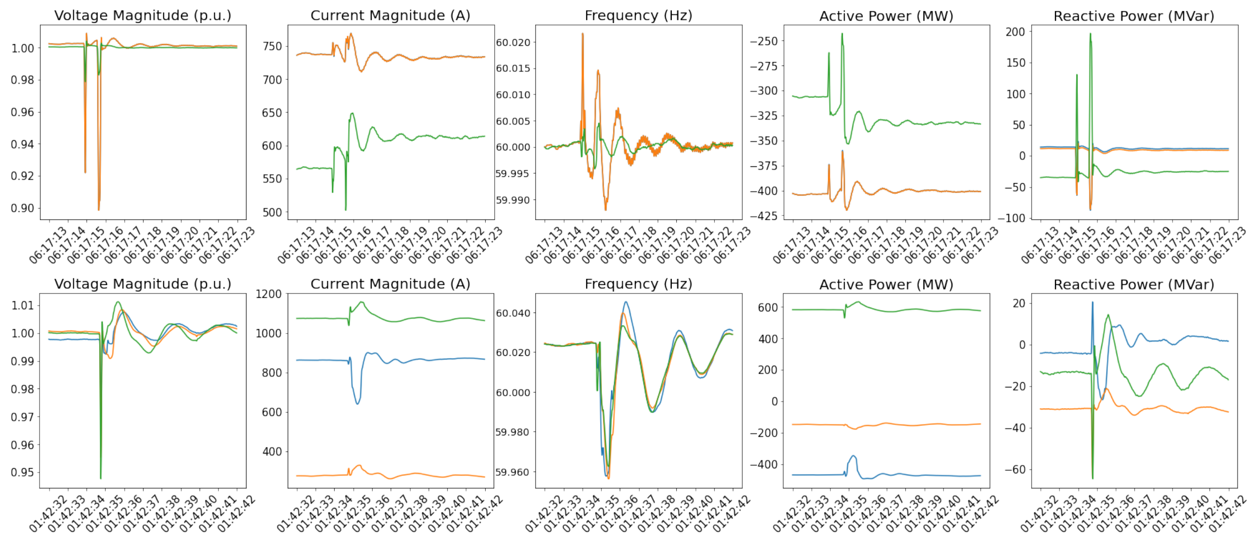

Figure 6. Ringdown oscillation examples. The plots show 10 seconds of data.

- Oscillation may occur.
- Frequency spikes.

Project 4 also noted that not all line trips were caused by faults. For instance, some trips were due to the activation of remedial action schemes (RAS). Similarly, not all faults resulted in trips. Faults that cleared normally were unlikely to be documented in utility logs. One may ascertain whether a detected fault caused a trip by checking whether current flow over at least one line decreased to zero, or whether step changes were seen in active/reactive power following the event. Detecting faults that did not trip any equipment may help identify the presence of high-impedance faults.

4.2 Other Events

Awardee projects also found evidence of other events such as line reconnections and auto-reclosure actions (an example found by Project 4 is shown in Fig. 9). Reliable identification of auto-reclosure actions can be beneficial to utilities. For example, if frequent auto-reclosures are reported from one location, it might indicate the presence of encroaching vegetation coming in frequent contact with power lines. If such locations are identified, utilities can take necessary preventive actions before any major faults happen.

The events and anomalies detected by the project teams also contained many events which do not match the signatures for any of the categories described in this chapter. Many of these disturbances were small and may not be of interest to utilities. Project 2 used an event label called 'operation' in their classification model to account for some such mundane step changes and switching actions present in the data.

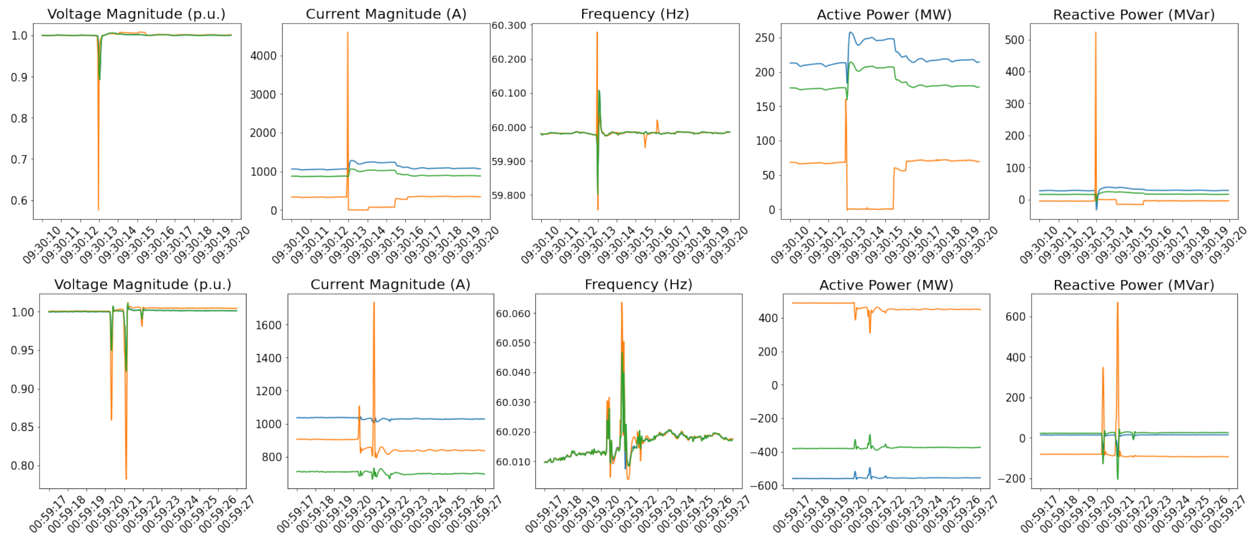

Figure 7. Line trip examples. The plots show 10 seconds of data.

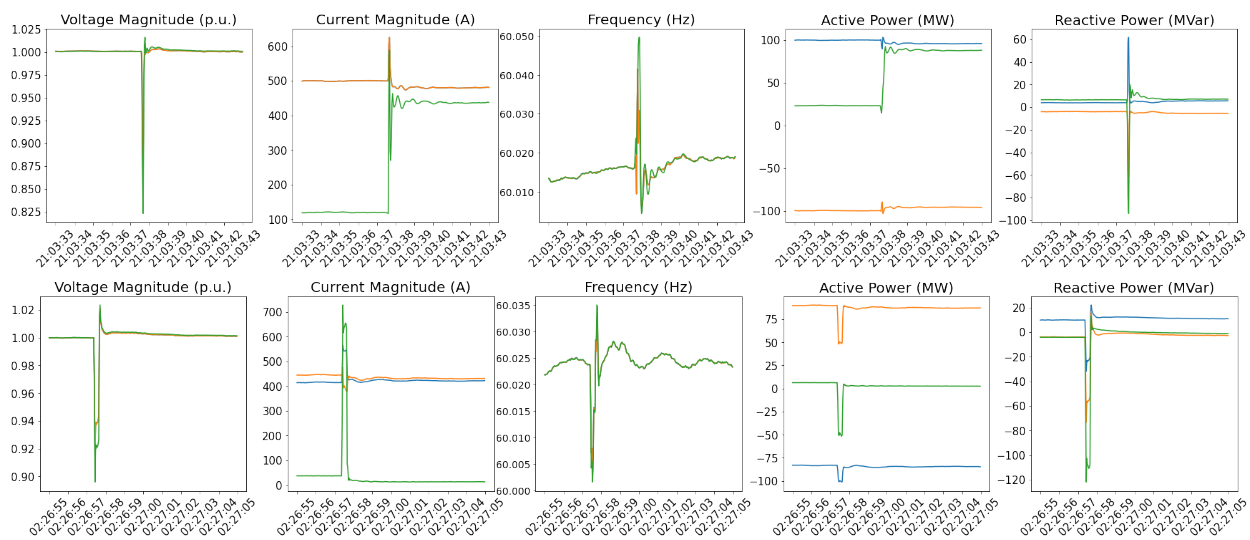

Figure 8. Transformer trip examples. The plots show 10 seconds of data.

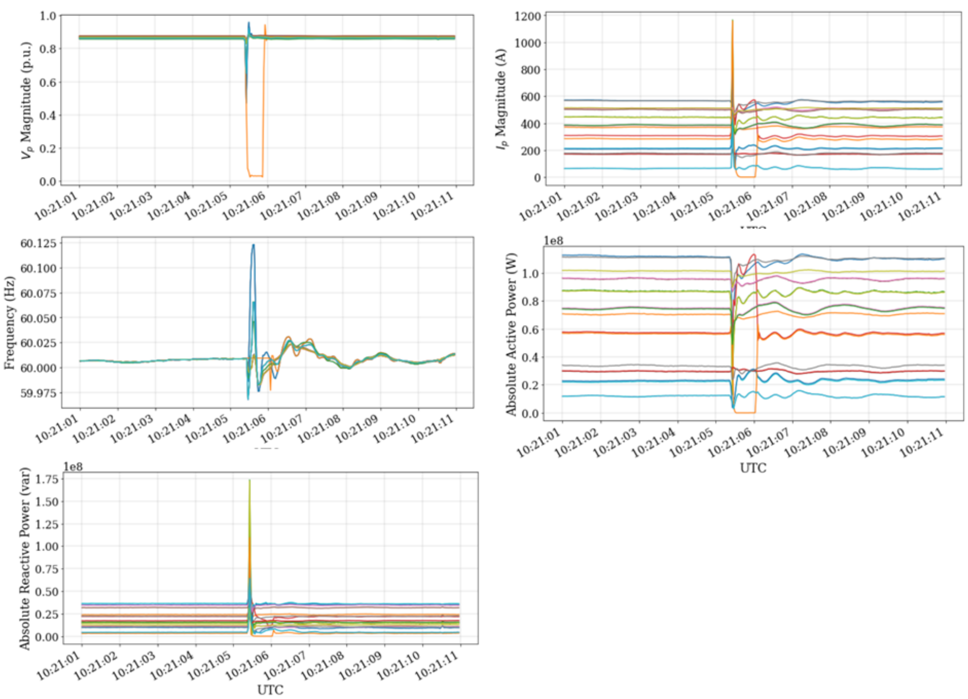

Figure 9. Example of auto-reclosure action identified by Project 4

5.0 Conclusion

The overarching aim of FOA 1861 was to accelerate the development of data-driven and AI/ML-based approaches for converting large volumes of PMU data into actionable insights. The awardees developed a variety of data repair, visualization, event detection and classification algorithms, as well as robust infrastructure for data handling and computation aimed at efficient execution of the end-to-end process from data retrieval to event classification. They extracted a rich repository of labeled events, and also investigated various methodologies for detecting precursors to impending failures, mitigating GPS spoofing attacks, and more.

While it has been recognized for a long time that ML models can aid grid operations and alleviate the cognitive load of operators, formulating effective solutions and translating theoretical solutions to commercial tools has been a challenge. There are several reasons behind this, including the lack of good quality labeled field data for training models, industry resistance to relying on opaque black-box algorithms in the context of critical infrastructure operations, and the lack of demonstrated reliability and generalizability of ML models. As data-driven models are not yet ready to provide full diagnosis of events and their root causes in the field, the immediate focus of real-world deployments should be on algorithms that help in automating aspects of operator and engineering workflows. Ideally, these algorithms would-

- better filter or highlight information, and
- augment operator memory, knowledge retrieval and keep track of the latest events.

To this end, statistical feature-based methods may be deployed in the near term to aid operators in reducing the amount of data they have to review. Further, these methods can add to the repository of properly-labeled events that can go on to train improved ML models.

5.1 Operationalizing the Insights

Based on insights gathered during the project period, the awardees put forth recommendations for future actions, some of which are listed below.

- Most awardees struggled with the inconsistent labeling of events in utility-provided logs. Many of them suggested fostering discussions in the working groups centered on formulating standardized labelling practices and using uniform terminology. A possible enhancement could include a hierarchical label definition, reflecting different degrees of knowledge about an event. Standard labeling practices across organizations will help create common datasets that can be used in benchmarking the performance of different algorithms.

- Awardees also struggled with data quality issues- a lot of the measurements had artifacts such as spurious drops, inconsistent unit usage and sudden spikes in values. As many of the awardees spent significant time looking at these issues, Project 4 suggests the creation of a catalogue of such signatures that can be shared with practitioners to further industry knowledge of common PMU data quality issues.

- Project 3 found that a high-performance HDF5 data storage format provided much faster computation performance and had much lower memory requirements than the parquet format that the FOA 1861 data was stored in. They recommend using the same schema for big data storage in any follow-on work.

- Algorithms developed and refined with the support of this FOA can be implemented as tools within existing commercial wide-area measurement system (WAMS) software platforms. Extracted event signatures will also help in refining the methodologies in the existing WAMS platforms. Project 4 showed an example of this by developing an application implementing their trained semi-supervised classification model on MindSphere, the industrial IoT solution offered by Siemens.

There are other low-hanging opportunities for operationalizing the FOA research outcomes as well. For example, the awardees developed algorithms that can reliably detect generator loss events. These classifiers can be utilized to automatically export interesting events to event analysis and model validation/calibration tools. Similarly, ringdown detectors can be set up to automatically choose events for modal analysis. The repository of event signatures can also be a valuable resource for developing educational tools for operators and university students.

5.2 Further Work

There are many avenues for future work building on the FOA 1861 effort. Some suggestions put forth by the awardees are stated below.

- *Big data visualization tools*: In the foreseeable future, ML will not be replacing human operators from control rooms. Rather, it will aid human decision-making. A critical part of creating algorithms that work in effective collaboration with humans is the design of effective interfaces and visualizations through which algorithm outputs can be communicated. In the electric grid context, visualization design is particularly challenging—and key—because large quantities of complex information must be conveyed to cognitively strained operators. Hence, future research on effective and intuitive visualization tools that reduce cognitive burden on human operators must be encouraged. (Project 1)

- *Creation of golden datasets*: Project 4 recommends the creation of 'golden' datasets that contain clearly labeled and validated (ideally following a standardized labeling standard) events. Such a dataset may be used to benchmark the performance of different proposed ML algorithms, and help users choose which approach is best-suited for their application.

  PNNL researchers are in the process of creating such a dataset using the huge repository of events extracted by the FOA awardees (one of the key outcomes of this FOA). All events in this signature library will have been validated by domain experts, and will be made publicly available. The synthetic dataset pmuBAGE created by Project 5 will also provide a common benchmark for algorithm performances.

- *Incorporating feedback learning*: As the lack of labeled data may inhibit the performance of ML models, Project 1 suggests exploring feedback learning techniques. As the name suggests, feedback learning enables the solicitation of user feedback for ML systems. Human users could flag or confirm ML predictions, allowing algorithms to learn in deployment. Given the paucity of labeled grid data, and the poor understanding of how training data generalizes across contexts, this learning on the fly could, over time, greatly improve ML algorithm performance.

6.0 References

[1] Department of Energy (DOE) Office of Electricity, "Funding Opportunity Announcement (FOA) Number: DE-FOA-0001861, Big Data Analysis of Synchrophasor Data, Amendment 000001," Oct. 2018.

[2] J. S. Banning, "FOA 1861 subject matter expert-data description," Tech. Rep. PNNL-SA-148400, Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2019.

[3] J. S. Banning, J. D. Follum, and E. S. Andersen, "FOA 1861 data curation overview," Tech. Rep. PNNL-32303, Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2021.

[4] S. Biswas, J. Follum, J. S. Banning, S. Kincic, and E. Andersen, "Anonymization effectiveness assessment in the FOA 1861 synchrophasor dataset," tech. rep., Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2022.

[5] K. D. Jones, A. Pal, and J. S. Thorp, "Methodology for performing synchrophasor data conditioning and validation," *IEEE Transactions on Power Systems*, vol. 30, no. 3, pp. 1121–1130, 2015.

Appendix A – List of Publications

A.1 Technical Products

A.1.1 Project 1

Publications

1. PingThings Analytics, "Final Scientific/Technical Report: Combinatorial Evaluation of Physical Feature Engineering, Classical Machine Learning, and Deep Learning Models for Synchrophasor Data at Scale," PingThings, CA, USA, Tech. Report., April 2022. Available: https://www.osti.gov/servlets/purl/1864556

A.1.2 Project 2

Publications

1. GE Research, "Final Scientific/Technical Report: Phasor-Measurement-Unit-Based Data Analytics Using Digital Twin and PhasorAnalytics Software", GE Research, Niskayuna, NY. Available: https://www.osti.gov/servlets /purl/1828164
2. V. S. Kumar, T. Wang, K. S. Aggour, P. Wang, P. J. Hart and W. Yan, "Big Data Analysis of Massive PMU Datasets: A Data Platform Perspective," in *2021 IEEE PES ISGT NA*, 2021, pp. 1-5.
3. T. Wang et al., ''Grid event signature identification through large, real-world PMU datasets,'' Gen. Electr. Res., Niskayuna, NY, USA, Tech. Rep. 2021092068790, 2022.
4. P. Hart et al., "Application of Big Data Analytics and Machine Learning to Large-Scale Synchrophasor Datasets: Evaluation of Dataset 'Machine Learning-Readiness'," in *IEEE Open Access Journal of Power and Energy*, vol. 9, pp. 386-397, 2022, doi: 10.1109/OAJPE.2022.3197553.

Invention Disclosures

1. ID89835397: "A computation-efficient feature generation engine for Phasor Measurement Unit data set", disclosed on 9/30/2020
2. ID89888522: "PMU data-based time domain ringdown detection from a hyperspatial perspective", disclosed on 2/3/2021

A.1.3 Project 3

Publications

1. Schweitzer Engineering Laboratories, Inc., "Final Scientific/Technical Report: Machine Learning Guided Operational Intelligence from Synchrophasors", SEL Inc., Pullman, WA, July 2021.

Available: https://www.osti.gov/ servlets/purl/1828371

2. Lassetter, E. Cotilla-Sanchez, and J. Kim, "Using Critical Slowing Down Features to Enhance Performance of Artificial Neural Networks for Time-Domain Power System Data", *International Conference on Smart Energy Grid Engineering*, 2021.

3. S. De Silva, J. Kim, and E. Cotilla-Sanchez, "Data Driven Sparse Error Correction for PMU Measurements under GPS Spoofing Attacks", in *IEEE PES ISGT NA*, pp.1-5, February 2021.

4. D. Senaratne, J. Kim, and E. Cotilla-Sanchez, "Spatio-Temporal Frequency Domain Analysis of PMU Data for Unsupervised Event Detection", in *IEEE PES ISGT NA*, pp.1-5, February 2021.

A.1.4 Project 4

Publications

1. Siemens Corporation, "FInal Scientific/Technical Report: MindSynchro", December 2021.

2. B. P. Leao et al. "Big Data Processing for Power Grid Event Detection", in *IEEE Big Data 2020*. Available: https://www.osti.gov/biblio/1764612

3. Y. Du et al. "Physics-Based Feature Extraction from Bulk Time-Series PMU Datasets for Event Detection", in IEEE PES General Meeting 2021. Available: https://www.osti.gov/biblio/1826178

4. T. Lan et al. "Unsupervised Power System Event Detection and Classification Using Unlabeled PMU Data", in *IEEE ISGT Europe 2021*. Available: https://www.osti.gov/biblio/1828376

5. B. P. Leao et al. "Unleashing the Power of Industrial Big Data through Scalable Manual Labeling", in *NeurIPS Worshop on Data-Centric AI 2021*. Available: https://datacentricai.org/neurips21/papers/84_CameraReady_MindSynchro_NeurIPS_Data_Centric_AIfinal.pdf

6. D. Fradkin and B. P. Leao. "Improving Smart Grid Data Quality via Anomaly Detection", in *IEEE SGSMA 2021* (submitted).

Invention Disclosures

1. "Method for detection of relevant events in power grid operation" (Siemens Docket No. 2020E05865 US)

2. "Method for monitoring and detection of relevant power grid events" (Siemens Docket No. 2020E11809 US)

3. "An Interface for Grid Event Annotation from PMU Data" (Siemens Docket No. 2020E13172 US)

4. "Data labeling and label validation through ensembles of clustering models" (Siemens Docket No. 2021E01748US)

5. "Multiple Instance Learning for Grid Event Detection from PMU Data" (Siemens Docket No. 2021E07661 US)

6. “Event driven topological estimation using correlated projections” (Siemens Docket No. 2021E10200 US)

A.1.5 Project 5

Publications

1. University of California Riverside, “Final Scientific/Technical Report: Discovery of Signatures, Anomalies, and Precursors in Synchrophasor Data with Matrix Profile and Deep Recurrent Neural Networks”.

2. B. Foggo and N. Yu, “Online PMU Missing Value Replacement via Event-Participation Decomposition,” in *IEEE Trans. on power systems*, Vol. 37, No. 1, pp. 488-496, 2021.

3. Y. Cheng, B. Foggo, K. Yamashita, and N. Yu, “Short-term Forecasting of PMU Data by Attentional Seq2Seq LSTM with Prior Knowledge Matrix and Magnitude Direction Coupling”.

4. J. Shi, B. Foggo, X. Kong, Y. Cheng, N. Yu, and K. Yamashita, “Online Event Detection in Synchrophasor Data with Graph Signal Processing,” *IEEE SmartGridComm*, 2020.

5. X. Kong, B. Foggo, K. Yamashita, N. Yu, “Online Voltage Event Detection Using Synchrophasor Data with Structured Sparsity-Inducing Norms,” in *IEEE Trans. on power systems*, vol. 37, no. 5, pp. 3506-3515, Sept. 2022, doi: 10.1109/TPWRS.2021.3134945.

6. Y. Cheng, N. Yu, B. Foggo and K. Yamashita, “Online Power System Event Detection via Bidirectional Generative Adversarial Networks,” in *IEEE Transactions on Power Systems*, Early Access, doi: 10.1109/TPWRS.2022.3153591.

7. J. Shi, B. Foggo, and N. Yu, “Power System Event Identification based on Deep Neural Network with Information Loading,” in *IEEE Trans. on power systems*, Vol. 36, No. 6, pp. 5622-5632, 2021.

8. Y. Cheng, N. Yu, and K. Yamashita, “Adversarial Attacks on Deep Neural Networks Based Power System Event identification,” *IEEE ISGT Aisa 2022*, under review, 2022.

9. Jie Shi, K. Yamashita, and N. Yu, “Power System Event Identification with Transfer Learning Using Large-scale Real-world Synchrophasor Data in the United States,” in *IEEE PES ISGT NA*, 2022.

10. K. Yamashita, B. Foggo, X. Kong, Y. Cheng, J. Shi and N. Yu, “A Dynamic Behavior-Based Bulk Power System Event Signature Library With Empirical Clustering,” in *IEEE Access*, vol. 10, pp. 96307-96321, 2022, doi: 10.1109/ACCESS.2022.3205321.

11. B. Foggo, K. Yamashita, and N. Yu, “pmuBAGE: The Benchmarking Assortment of Generated PMU Data for Power System Events -Part I: Overview and Results,” in *IEEE Trans. on power systems*, under review, 2022. arXiv:2204.01095v1

12. B. Foggo, K. Yamashita, and N. Yu, “pmuBAGE: The Benchmarking Assortment of Generated PMU Data for Power System Events -Part II: Methodology and Model Details,” in *IEEE Trans. on power systems*, under review, 2022.

13. X. Kong, K. Yamashita, B. Foggo, and N. Yu, “Dynamic Parameter Estimation with Physics-based Neural Ordinary Differential Equations,” in *IEEE PES General Meeting*, July, 2022, Denver, Colorado.

A.1.6 Project 6

Publications

1. University of Nevada Reno, “Final Scientific/Technical Report: A Robust Event Diagnostics Platform: Integrating Tensor Analytics and Machine Learning into Real-time Grid Monitoring”.

2. A. Ghasemkhani, A. Darvishi, I. Niazazari, A. Darvishi, H. Livani, and L. Yang, “Deepgrid: Robust deep reinforcement learning-based contingency management,” in *IEEE PES ISGT NA*, February 2020.

3. I. Niazazari and H. Livani, “Attack on Grid Event Cause Analysis: An Adversarial Machine Learning Approach,” in *IEEE PES ISGT NA*, February 2020.

4. P. Zhao, P. Chen, S. Wang, and X. Lin, “Towards Query-Efficient Black-Box Adversary with Zeroth-Order Natural Gradient Descent,” *AAAI*, Feburary 2020.

5. M. Cheng, S. Singh, P. Chen, S. Liu, and C. Hsieh. ”Sign-OPT: A Query-Efficient Hard-label Adversarial Attack.” in *ICLR* 2020.

6. A. Ghasemkhani, I. Niazazari, Y. Liu, H. Livani, V. A. Centeno and L. Yang, “A Regularized Tensor Completion Approach for PMU Data Recovery,” in *IEEE Trans. on Smart Grid*, doi: 10.1109/TSG.2020.3030566.

7. T. Idé, A. Dhurandhar, J. Navratil, M. Singh and N. Abe, “Anomaly Attribution with Likelihood Compensation,” in *AAAI 21*, February 2021.

8. I. Niazazari, H. Livani, A. Ghasemkhani, Y. Liu, and L. Yang, “Event cause analysis in distribution networks using synchro waveform measurements,” in *North American Power Symposium (NAPS)*, April 2021.

9. I. Niazazari, Y. Liu, A. Ghasemkhani, S. Biswas, H. Livani, L. Yang, and V. A. Centeno, “PMU-data-driven event classification in power transmission grids,” in *IEEE PES ISGT NA*, February 2021.

10. A. Ghasemkhani, Y. Liu, and L. Yang, “Low-rank tensor completion for PMU data recovery,” in *IEEE PES ISGT NA*, February 2021.

11. T. Ide, G. Kollias, D.T. Phan, N. Abe, ”Cardinality-Regularized HawkesGranger Model”, in *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, 2021.

12. Y. Liu, L. Yang, A. Ghasemkhani, H. Livani, V. A. Centeno, P.-Y. Chen, and J. Zhang, “Robust event classification using imperfect real-world PMU data,” arXiv preprint arXiv:2110.10128, 2021.

13. A. Ghasemkhani, Y. Liu, and L. Yang, “Real-time event detection using rank signatures of real-world PMU data,” in *IEEE PES General Meeting*, Denver, CO, USA, July, 2022.

14. Y. Liu and L. Yang, “Weakly supervised event classification using imperfect real-world PMU data with scarce labels,” in *IEEE PES General Meeting,* Denver, CO, USA, July, 2022.

A.1.7 Project 7

Publications

1. Iowa State University of Science and Technology, “Final Scientific/Technical Report: Robust Learning of Dynamic Interactions for Enhancing Power System Resilience”.

2. Y. Yuan, Y. Guo, K. Dehghanpour, Z. Wang and Y. Wang, ”Learning-Based Real-Time Event Identification Using Rich Real PMU Data,” in *IEEE Trans. on Power Systems*, vol. 36, no. 6, pp. 5044-5055, Nov. 2021, doi: 10.1109/TPWRS.2021.3081608.

3. Y. Yuan, Z. Wang and Y. Wang, ”Learning Latent Interactions for Event Identification via Graph Neural Networks and PMU Data,” in *IEEE Trans. on Power Systems*, doi: 10.1109/TPWRS.2022.3158248.

4. C. Shang, J. Chen and J. Bi, “Discrete Graph Structure Learning for Forecasting Multiple Time Series” in *9th International Conference on Learning Representations (ICLR 2021)*, May 2021. Available: https://arxiv.org/abs/2101.06861

5. E. Dai and J. Chen, “Graph-Augmented Normalizing Flows for Anomaly Detection of Multiple Time Series”, in *10th International Conference on Learning Representations(ICLR 2022)*, April 2022. arXiv:2202.07857v2

6. “Federated Inference through Disambiguating Local Representations and Learning a Consensus Graph”, in preparation.

7. “A Safe Semi-supervised Learning Framework for Event Identification Using PMU Data” submitted to *IEEE Trans. on Power Systems*.

A.1.8 Project 8

Publications

1. Texas A&M Engineering Experiment Station,“Final Scientific/Technical Report: Big Data Synchrophasor Monitoring and Analytics for Resiliency Tracking (BDSMART)”.

2. M. Alqudah, M. Pavlovski, T. Dokic, M. Kezunovic, Y. Hu, Z. Obradovic, “Convolution-based Fault Detection Utilizing Timeseries Synchrophasor Data from Phasor Measurement Units,” *IEEE Trans. on Power Systems*, 2022. DOI: 10.1109/TPWRS.2021.3135336.

3. M. Pavlovski, M. Alqudah, T. Dokic, A. Abdel Hai, M. Kezunovic, Z. Obradovic, “Hierarchical Convolutional Neural Networks for Event Classification on PMU Data,” *IEEE Trans. on Instrumentation and Measurement*, Vol. 70, pp. 1-13, no. 2514813, 2021. DOI: 10.1109/TIM.2021.3115583.

4. A. Abdel Hai, T. Dokic, M. Pavlovski, T. Mohamed, D. Saranovic, M. Alqudah, M. Kezunovic, Z. Obradovic, “Transfer Learning for Event Detection from PMU Measurements with Scarce

Labels" *IEEE Access*, Vol. 9, pp. 127420 – 127432, September 2021. DOI: 10.1109/ACCESS.2021.3111727.

5. M. Kezunovic, Z. Obradovic, Y. Hu "Automated System-wide Event Detection and Classification Using Machine Learning on Synchrophasor Data," *CIGRE General Session*, Paris, France, Aug. 2022.

6. M. Kezunovic, Z. Obradovic, Y. Hu "Use of Machine Learning on PMU Data for Transmission System Fault Analysis," *CIGRE General Session*, Paris, France, Aug. 2022.

7. Z. Cheng, Y. Hu, Z. Obradovic, M. Kezunovic, "Using Synchrophasor Status Word as Data Quality Indicator: What to Expect in the Field?", in *IEEE Smart Grid Synchronized Measurement and Analytics Conference (SGSMA) 2022*, Split, Croatia, May 2022.

8. T. Dokic, R. Baembitov, A. Abdel Hai, Z. Cheng, Y. Hu, M. Kezunovic, Z. Obradovic, "A Single-Feature Machine Learning Method for Detecting Multiple Types of Events from PMU Data," in *IEEE Smart Grid Synchronized Measurement and Analytics Conference (SGSMA) 2022*, Split, Croatia, May 2022.

9. H. Otudi, T. Dokic, T. Mohamed, M. Kezunovic, Y. Hu, Z. Obradovic, "Line Faults Classification Using Machine Learning on Three Phase Voltages Extracted from Large Dataset of PMU Measurements," in *HICSS-55 Conference*, Hawaii, USA, January 2022.

10. R. Baembitov, T. Dokic, M. Kezunovic and Z. Obradovic, "Fast Extraction and Characterization of Fundamental Frequency Events from a Large PMU Dataset Using Big Data Analytics," in *HICSS-54 Conference*, Hawaii, USA, January 2021.

**Pacific Northwest
National Laboratory**

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

www.pnnl.gov