Paper ID – 224 Session 2022
C2: Power System Operation&Control
PS1 - System Control Room Preparedness: Today and in The Future

Automated System-wide Event Detection and Classification Using Machine Learning on Synchrophasor Data

Mladen KEZUNOVIC* — Texas A&M University, U.S.A. — m-kezunovic@tamu.edu

Zoran OBRADOVIC — Temple University, U.S.A — zobrad@gmail.com

Yi HU — Quanta Technology, U.S.A. — YHu@Quanta-Technology.com

SUMMARY

As the number of phasor measurement units (PMUs) deployed in a power system increases, and their data volume streamed to the control canter intensifies, operators are facing challenges related to the analysis of such data, which need to be observed and responded to as the measurements are displayed in the Control Room. Humans are generally unable to process such large amount of data efficiently and rapidly. There is an apparent need for automated ways to analyze the data, extract actionable information about occurrence of specific events, and characterize the events quickly and cost effectively. This paper discusses the use of machine learning (ML) to facilitate such tasks by providing automated, highly computationally efficient, and cost-effective ways of extracting actionable information from synchrophasor big data in real-time.

We developed Big Data Smart (BDSmart) ML-based prototype tool for the Control Room use that automatically analyses data properties from synchrophasor system measurements taken across the three grid Interconnections in the USA (Western, Eastern and ERCOT). The data collected from several hundreds of PMUs located across the Interconnections over a period of two years have been made available for our extensive study. As a result, we were able to identify a number of big data properties that influence how ML methodology is applied to select, develop, train and test the data models that can eventually be used for the tool implementation. The resulting set of candidate algorithms spans unsupervised, supervised, semi-supervised and transfer-learning approaches. Many ML techniques, such as decision trees, multinomial logistic regression, feed-forward neural networks, K-nearest neighbor, multiclass support vector machine, and single and multi-channel convolutional neural networks, are implemented, and their performance is examined. We offer the results from testing the data models.

The novelty of our study is in the approaches for bad data detection and mitigation, selection of a simplified feature for event detection, and data label improvements. As a result, we came up with a list of recommendations for the utilities on how to improve the PMU recording practices to cater to the future ML applications aimed at automating the analysis of synchrophasor data.

KEYWORDS

Synchrophasors, Machine Learning, System-wide Event Analysis, Automated Data Analytics

1. INTRODUCTION

To develop ML models, an abundance of historical data is needed. Due to the electric grid security and various commercial reasons, the synchrophasor data shared with the ML model developers may have to be restricted. We discuss how the data limitations may affect the model development using a use case where the data providers have decided to impose rather conservative approach in sharing the data and information about the synchrophasor system. This creates several constraints on how the data models may be selected, structured and eventually trained and tested.

The Critical Electric Infrastructure Information (CEII)

According to the U.S. Department of Energy, "CEII is a category of controlled unclassified information about a system or asset of the bulk-power system, whether physical or virtual, that if destroyed or incapacitated, would negatively affect the United States' national security, economic security, public health or safety, or any combination of such effects. A CEII designation exempts the information about physical and virtual assets of the bulk-power system from public release under the Freedom of Information Act and other laws requiring government disclosure of certain information or records. As a general principle, DOE will not designate information as CEII if it has been made publicly available by an owner or generator of the CEII previously" [1]. As a result of such requirements, we received no information about either the grid topology or the PMU placements. We did receive the event labels that were produced by the data owners, but it was not known which part of the Interconnection was involved in the system-wide event; hence it was not clear whether the labels were accurate.

The Data Sharing Between Transmission Operators and Independent System Operators

In addition to CEII restriction in data sharing, the Interconnections have their own Data Sharing Agreements that may affect the information released to the ML Model developers [2]. Such agreements require close coordination among the data owners before the covered data is released. In our case, such coordination resulted in a decision to provide no data about the synchrophasor system design, including PMU make or vintage, time-synchronization or time-stamping approaches, the communication system properties, or configuration parameters for PMUs or Phasor Data Concentrators (PDC). This significantly affects our ability to fully understand the causes for frequent occurrence of bad data issues such as time synchronization, missing data, or unreasonable data.

The Inability to Fully Describe System-wide Events Across Interconnections

The focus of our analysis was on the system-wide events such as fundamental-frequency deviations that may lead to instability issues and low frequency oscillations among groups of generators. According to the Electric Information Agency, "the US electric system is made up of three Interconnections and 66 Balancing Authorities" [3]:

- The Eastern Interconnection encompasses the area east of the Rocky Mountains and a portion of northern Texas. The Eastern Interconnection consists of 36 balancing authorities: 31 in the United States and 5 in Canada.
- The Western Interconnection encompasses the area from the Rockies west and consists of 37 balancing authorities: 34 in the United States, 2 in Canada, and 1 in Mexico.
- The Electric Reliability Council of Texas (ERCOT) covers most, but not all, of Texas and consists of a single balancing authority.

In such a setting of dispersed authorities, an explanation of a system-wide event may require coordination of multiple entities, and it was not known to the ML data model developers how widespread a system event was and who was involved in developing the labels describing the event type and time of occurrence. This resulted in an uneven distribution of label quality and perhaps in an inaccuracy of the event labels accompanying data recordings for different Interconnections.

2. PROPERTIES OF SYNCHROPHASOR FIELD MEASUREMENTS

To illustrate some of the most common recorded measurement issues, we focus on some generic data properties that have major impact on the development of ML data models. We illustrate each of the generic issues with some examples from the datasets we used for our study, which as mentioned before comprises data from several hundreds of PMUs recorded over a period of two years from all three US Interconnections.

Bad Data Detection and Mitigation

A scan of the provided datasets for three Interconnections has revealed several major data quality issues such as: (1) *Missing data*: the amount of data missing range from ~2% to ~60% on an individual PMU basis; and (2) *Unreasonable data*: this includes outliers (e.g. voltage magnitude has negative values or well beyond twice the nominal values, phase angle values outside ± 180 degree range) and non-measurement data (e.g. flat 60 Hz value for frequency measurement in some PMUs instead of a constantly changing frequency value in other PMUs). The datasets included the bit mapped 2-Byte STAT word where each bit is defined by the IEEE standard C37.118.2-2011 to indicate the data quality status for each PMU at each time stamp [4]. However, after an examination of these STAT words we concluded that they cannot be used as a reliable data quality indicator for the provided data. The main issues are inconsistencies in setting these bits among PMU vendors and users, no information of the version of the standard implemented by vendors, and no information for the synchrophasor systemin which some system components such as PDCs could also alter the bits in STAT words.

This study takes a different approach in mitigating the bad data quality issues. No comprehensive data cleansing and reconstruction effort was made upfront to the datasets in this project. Instead, an assessment is performed to determine whether a particular data quality issue may have any impact on system-wide event detection by a specific ML model. Appropriate mitigation actions are taken if there is an impact. For system-wide frequency event detection, some PMUs missing data during such event would not prevent the event to be detected, so no mitigation action for this issue is taken. Flat 60Hz data or unreasonable outlier values of the measured frequency impacts the system-wide frequency event detection, and therefore those values are excluded from being used by the ML models.

Curation of Imprecise or Missing Labels through Visual Inspection

Event logs for about 75% of the data in the datasets (i.e., for every 8-week data, event log for first six-weeks) is provided by the data providers. However, the categorization and labelling of various types of events are not consistent for all three Interconnections. As an example, in the Eastern and ERCOT Interconnections, a generator tripping event is logged under the category of "Generator" and the cause of the event as "Trip", while the Western Interconnection does not have the "Generator" category and the "Trip" cause under it. The generator tripping event in the Western Interconnection is likely logged under the category of "Frequency" events with no causes identified, because loss of a large amount of supply due to generator tripping typically will lead to a frequency deviation event (Figure 1 (a)). However, the same "Frequency" category also includes other types of frequency events such as inter-area oscillation events (Figure 1 (b)).

Inconsistency in categorizing and labelling events directly impacts the performance of ML models being trained. We have taken the time to visually inspect the logged events based on the information provided by the event logs, and relabeled some of the events in a consistent way. Training ML models using event data with cleaned labels has resulted in the performance improvement of the trained ML models [5].

Our visual inspection has also identified additional events that are not among the logged events. These additional events may have occurred outside the control areas of the data providers, so even though the events might have been "seen" by PMUs of the data providers, these events might have not been logged by the data providers.

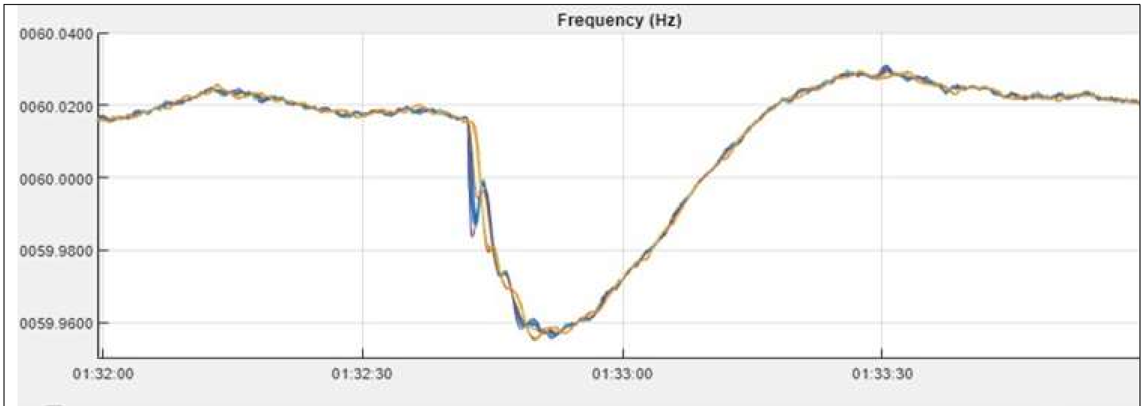

(a) "Frequency" event 1 – Frequency deviation

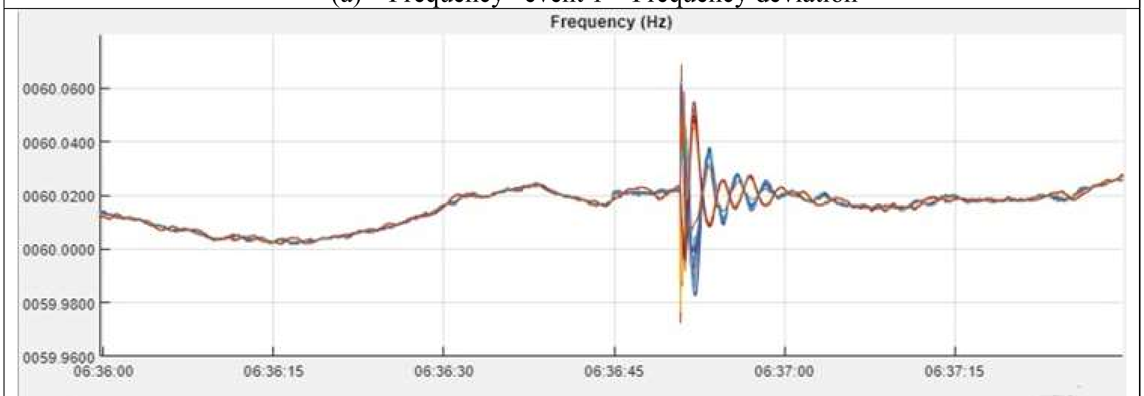

(b) "Frequency" event 2 – Inter-area oscillation

Figure 1: Two logged "Frequency" events seen by PMUs in the Western Interconnection

3. SYSTEM-WIDE EVENTS OF INTEREST

Inspection of the data labels indicates several categories of events that are captured by PMUs over an extended period of time. Closer analysis identifies events that may have been attributed to the maintenance procedures. So, we decided to develop the ML models that will only address the operational events that may affect the system performance, which is the responsibility of the system operators. This has imposed the requirement that our ML models are computationally efficient so that the results can be produced timely for the operators to be able to react with possible mitigation measures.

Fundamental Frequency Deviations

The fundamental frequency of an interconnected power system constantly varies as a result of the real-time imbalance between the demand and supply. The frequency goes well below (or above) the nominal system fundamental frequency such as 60Hz in US when the total system demand is more (or less) than the total supply. As the total demand during a day is constantly changing, and now adding the supply fluctuations due to high-level penetration of variable energy resources, the fundamental frequency is fluctuating around the nominal value most of the time as a result of the constant supply-demand imbalance and the re-balancing actions by system operators. Typically, such frequency deviations are controlled by the system operators to be within a targeted range (e.g., ± 0.1Hz from the nominal fundamental system frequency) and deviations within this range are considered normal.

Fundamental frequency deviations from the nominal value may exceed this controlled range for certain frequency events, such as low frequency oscillations and large frequency excursions. Frequency events generally can be "seen" by all PMUs within an interconnection as shown in Figure 1 (a) and (b).

Low Frequency Oscillations

Various types of low frequency oscillations occur in an interconnected power system all the time. Most of the oscillations are short-lived natural oscillations that have little impact on power system normal operation and reliability. However, certain types of low frequency oscillations, namely an inter-area oscillation (i.e., two groups of generators in two different areas oscillate with each other) or a forced oscillation driven by a source of oscillation such as a connection of a group of renewable generators, could impact a wide area of an interconnected system and may lead to system instability or severely limit the power transfer capacity of transmission lines if not timely mitigated [6].

Sustained low frequency oscillations are longer duration events that protective relays and digital fault recorders are unable to capture due to lack of triggering conditions and/or fixed short event record length for each event. Low frequency oscillations are also not visible from the low-resolution SCADA RTU data. Streaming PMU data are the most suitable type of data for detecting those long duration low frequency inter-area and forced oscillations, and for locating the source of the forced oscillations.

Generator or Transmission Line Tripping

Generator(s) tripping causes sudden large supply shortage in the area where generator(s) was(were) connected, which could lead to large frequency drop in the area beyond the controlled range. Major tie-line(s) tripping could create two areas: one with excessive supply and one with large supply shortage leading to frequency jump in the excessive supply area and frequency drop in the other. As a consequence, the entire Interconnected system may be impacted by a generator tripping or a transmission tie-line tripping. The maximum frequency jump or drop depends on the total demand and remaining supply in each area immediately after such line tripping occurs. If the frequency drop exceeds the low frequency load shedding thresholds, load shedding will occur affecting large number of customers. When the frequency jumps above the generator's high frequency tripping thresholds, additional generators may be tripped off, which may lead to cascading system collapse.

PMU data could be used to help identify if there is any triggering event (e.g., a fault) that may have led to the generator or line tripping event in the first place, and to predict the maximum drop or jump long before they occur. Such pre-cursors could be used by the system operators in their real-time decision making or by automated protection and control systems to take pre-emptive actions.

4. ML PIPELINE FOR EVENT DETECTION AND CLASSIFICATION

A sequence of signal transformations and modelling steps is needed to enable implementation of ML for system-wide event analysis from PMU measurements. Towards this objective, we developed a variety of feature engineering approaches and models, which are briefly summarized below. We describe knowledge based automated methods for selecting explanatory variables associated with event labels and discuss how such training data is used for deep learning. In a complementary approach, when representative training data was not available, we used a transfer learning approach to select relevant labelled data and propagate experience to develop models where the available data was insufficient.

Knowledge Based and Automated Feature Engineering

In ML classification systems, a training set of examples is used to learn a function that maps given inputs to the outputs. Each input is a vector of explanatory variables describing a certain event, while the output is the type of the event. In system-wide analysis the problem is temporal, and measurements are obtained from multiple PMUs collocated with event log. Therefore the problem is formulated such that given a signal segment $\mathbf{s}(t-\Delta, t+\Delta) = [\mathbf{s}^{(1)}(t-\Delta, t+\Delta), \dots, \mathbf{s}^{(M)}(t-\Delta, t+\Delta)]$, from multiple anonymized PMUs (removed grid topology) the objective is to predict event type $y \in \{0, \dots, C\}$ that occurred at time $[t-\Delta, t+\Delta]$ by learning from scarce observations and low precision labels.

A very important modelling decision in the ML pipeline is signal transformation and selection of a vector of explanatory variables, called features, that are used for the model training and type prediction. This feature engineering process involves constructing a single vector of features corresponding to a single segment $\mathbf{s}(t-\Delta, t+\Delta)$, which can be based on elaborate domain knowledge or automated.

Traditionally, features for ML modelling are engineered based on expert-knowledge. Such an approach is usually easy to explain, it is computationally efficient and easy to deploy, but is typically limited to a specific application and might require additional data cleansing. In contrast, the aim of automated feature engineering approaches it to learn feature representation directly from data. Such an approach is usually computationally more demanding, but the benefits include internal filtering, data-driven dimensionality reduction, reduced need for domain expert involvement in model development, and less human error due to possible bias and knowledge leakage, potentially resulting in improved generalization.

For system-wide event detection from multiple PMUs, we have developed a knowledge-based feature engineering approach where features were extracted per PMU for each time segment $\mathbf{s}(t-\Delta, t+\Delta)$ by computing $(f_{max}-f_{min}) * (V_{max}-V_{min})$, where $f_{\max}$ and $f_{\min}$ are the maximum and minimum frequency values, and $V_{\max}$ and $V_{\min}$ are the maximum and minimum positive sequence voltage magnitude recorded by the selected *PMU* device, inside $(t-\Delta, t+\Delta)$ time window. In experiments conducted on data from the U.S. Western interconnection recorded at 38 PMUs over two years we show that a vector of such features is sufficient to capture whether an event has occurred within a time window, but such features are not suitable for classifying event types [7].

As an alternative, in the automated feature engineering approach, measurements observed at a specific PMU for a specific segment $\mathbf{s}(t-\Delta, t+\Delta)$ are *aggregated* to a lower resolution signal by binning the corresponding data into non-overlapping shorter buckets, where each bucket is summarized by the range of the raw measurements. This aggregation preserves the peaks as useful temporal properties for detecting and classifying events. Further aggregation is achieved by reducing aggregated information from all PMUs to a single sequence representing barycenter that preserves dynamic distances over all possible alignments between time sequences corresponding to reduced signal observed at individual PMUs. In this process, achieved through so called Soft Dynamic Time Warping, importance weights are assigned to information representing specific PMUs, and these weights are updated dynamically for each time interval [8]. This computationally demanding process results in low dimensional problem representation that indirectly filters noise through soft aggregation of the segments and learns features that are not biased by the most common signal type (corresponding to normal operation).

Deep Learning Event Detection and Classification Models from PMU Data

To leverage measurements from multiple PMUs for system-wide event detection and classification, we developed several deep learning-based ML methods, where mappings from the explanatory variables to the predicted classes were stacked in a hierarchy of increasing complexity and abstraction. Our approach is based on supervised learning and uses Convolutional Neural Network (CNN) models. In such an approach, an automated feature engineering technique discussed in the previous section is applied to measurements, and then parallel and concatenation-based convolutional neural networks that include multiple layers are utilized for event detection and classification [8]. Assuming univariate measurement segments (e.g., voltage magnitude) as inputs, we leveraged Single-Channel CNN (SC-CNN) as illustrated in Figure 2. In such an approach, each convolution layer transforms an input segment by convolving it with a certain filter and the result is passed through another filter, where filters are automatically learned. In this process, features produced by each filter are summarized through global average pooling, resulting in down sampled filtered features that are to a certain degree position and scale invariant. The pooled features are further mapped to a hidden representation used to detect an event and estimate class probabilities.

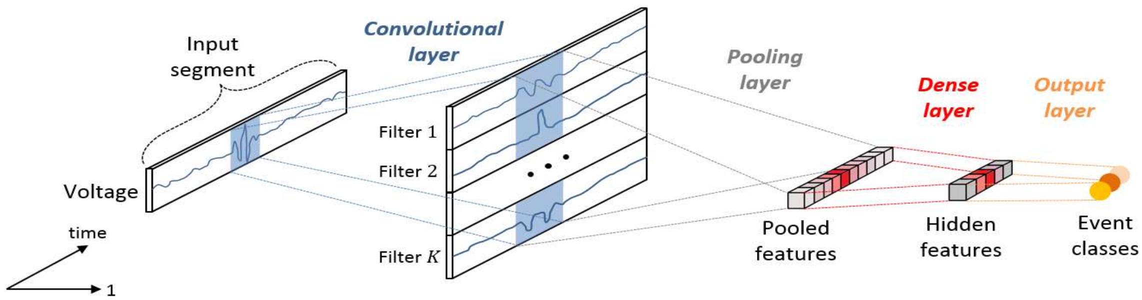

Figure 2: Single-Channel (SC) CNN utilizing voltage signal segments

We also considered two kinds of Multi-Channel CNNs. In Parallel Channel Filtering CNN (PCF-CNN) pooled features for voltage, current and frequency are produced separately, and the results are concatenated in a joint layer (Figure 3), while in Simultaneous Channel Filtering CNN (SCF-CNN) all 3 channels were mapped to a joint pooled feature as illustrated at Figure 4.

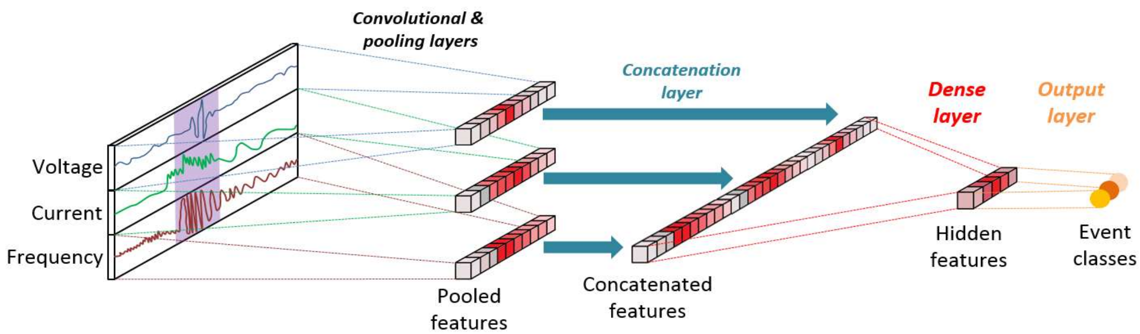

Figure 3: Parallel-Channel Filtering (PCF CNN) utilizing voltage signal segments

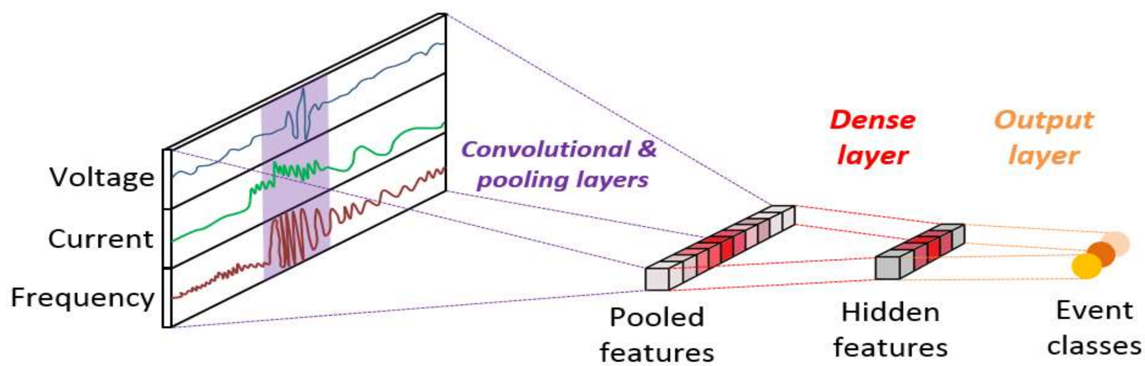

Figure 4: Simultaneous Channel Filtering (SCF CNN) utilizing voltage signal segments

We showed that CNN ML analysis was more accurate than traditional models in distinguishing frequency events from line faults and normal operation based on 2-years of PMU recordings at Western U.S. Interconnection. Using traditional models, we compared our results when we included Decision Tree [9], Logistic Regression [10], Multilayer Perceptron [11], and Multiclass Support Vector Machine [12]. The best predictions were consistently obtained using multi-channel CNNs [8].

Transfer Learning Models for Event Detection from PMU Data

Event detection and classification by supervised learning based on PMU data is a challenging task due to limited and imprecise event logs assigned to scarcely positioned PMUs in a large spatiotemporal system. An alternative of using unsupervised approaches might sound appealing since such an approach avoids the time-consuming and costly step of assigning labels manually. However, unsupervised detectors are usually less accurate as they do not benefit from labelled data that provide the possibility of correcting errors made by unsupervised detectors.

To overcome limitations of supervised and unsupervised ML, we utilized a transfer learning approach aimed at reducing the need for additional data labeling by leveraging some of labeled event data available from a related event detection source task and a smaller number of cases from a target task where we would like to detect events. Challenges of such an approach include violation of the following traditional machine learning assumptions: (1) the number of PMUs corresponding to dimensionality of the feature space of the source and target domain might be different; (2) marginal distributions of events could differ (covariate shift); and (3) the same behavior might have a different meaning in two domains (concept shift).

In our study [7], the first of these challenges is addressed by using an unsupervised Neural Network as an autoencoder to reduce dimensionality of the source and target domain (e.g., different interconnections) to the same latent space. Covariate shift and concept shift challenges were addressed by localized instance selection. This process transfers an annotated instance x_t from source domain D_S to target domain D_T if local structure of x_t is similar in D_S and D_T. We measure the distance between the centroids of the nearest neighborhood N_1 in D_S and N_2 in D_T, and the relative distance between the covariance matrices of N_1 and N_2. These distances are then used to learn the instance-transfer function by training SVM model on the target distribution.

We apply this approach to 2 years of field recordings collected by 38 PMUs in the U.S. Western Interconnection. The year-1 data is used for training, and concept shift was present between year-1 and year-2 events distribution as determined by a statistical test. Experiments conducted by relying on 20 to 700 characteristic events provide evidence that the developed transfer learning method significantly improves automated event detection based on PMU measurements when extensive labeling is costly or impossible to obtain. When compared to alternative state-of-the-art machine learning algorithms (unsupervised, semi-supervised, and supervised), transfer learning is significantly more accurate when learning to detect events from limited labeled observations (by transferring knowledge from only 20 representative labeled data instances) [7].

In addition, we tested if labeled events data observed through 38 PMUs placed in the U.S. Western Interconnection grid can be used to transfer relevant knowledge to the recordings from 178 PMUs located in the U.S. Eastern Interconnection. We questioned how the number of events transferred from the Western affect event detection accuracy in the Eastern Interconnection. We then evaluated our approach on one-year field-recorded data from the U.S.A. Eastern Interconnection. When the transfer learning method trained on 20 to 726 events from the Western Interconnection enriched with event records from the Eastern Interconnection observed in a previous year, transfer learning outperformed alternative state-of- the-art conventional ML methods. Accuracy improvements were particularly evident when utilizing a small amount of labeled event data. Therefore, we conclude that transfer learning could be a method of choice when labels are scarce.

5. LESSONS LEARNED

While the data sets that we studied have many constraints imposed by the data owners as described in the previous sections, we were able to learn many lessons that can help data owners and ML model developers to achieve further improvements in synchrophasor recording practices and model development methodologies. We share the observations below by categorizing them by the steps to be taken when the data models are developed, trained and tested and making recommendations associated with each of the steps.

Data Ingestion, Cleansing and Management

Field PMU data could experience a variety of data quality issues for various reasons. A full scan of the historical data to identify these issues and to understand the overall data quality situation is a vital step before using the data for ML model training to recognize system-wide types of events. It may not be necessary to always perform a comprehensive data cleansing before using the data if a specific data quality issue does not have any significant impact to the targeted type of event detection by a particular type of ML models. Data cleansing should only be performed for those data quality issues that have considerable impacts.

An extremely important data management function is data labelling. To label an event, the best available sources of information should be used. Using SCADA records to label events may create significant errors, particularly when determining the time occurrence of the event start/end. We recommend that the end users deploy time stamping available from PMUs to identify the start/end of an event rather than using rather inaccurate SCADA time reference.

Feature Engineering

The ability of machine learning techniques to detect, classify and anticipate events critically depends on selection of explanatory variables that are used for the model training and event type prediction. The quality of the monitoring tool based on multiple PMUs is significantly affected by (1) data anonymization (removal of the grid topology); (2) scarcity of observations; and (3) precision of event labels (log file).

Knowledge-based feature engineering in big PMU data is typically easy to explain. While it might be too specific and possibly biased, it is computationally efficient. When data is large and heterogenous, our results provide evidence that event detection and classification could be significantly improved by automated feature learning.

We recommend that careful selection of features be performed to determine which features may describe the events that are being detected. Simplifying the feature engineering process may save time, and subsequently the implementation cost but the ML models may produce unreliable and inconsistent results.

Model Development

When detecting and classifying events from many PMUs in big data, deep learning ML methods in general outperform traditional supervised learning models. In our experiments hierarchical deep-learning models consistently outperformed the standard multiclass variants. For event analysis voltage was more relevant than current and frequency, but the best results were achieved when signal from all 3 channels (phases) was analyzed jointly. When detecting events from limited even data, we found that transfer learning outperformed unsupervised, semi-supervised, and fully-supervised algorithms.

We recommend that PMU data to train the ML models be properly matched to the purpose of the model use. The three-phase data is preferred over positive sequence recording whenever feasible. If system-wide events can be accurately reproduced using physics-based models of an actual system we highly recommend that elaborate simulations of the events be performed to enhance the understanding of the event features, and eventually enhance training dataset for the ML models. The type of the ML model should be carefully selected to make the detection/classification processing computationally cost-effective, particularly if the ML model are applied to the streaming PMU data.

6. CONCLUSIONS

In conclusion, several study findings are worth mentioning:

- If the models are going to be used for mission-critical applications, then the best possible data should be provided to the model developers;
- Assigning imprecise or inaccurate labels to the data may significantly impair ML model performance, so special care should be given to improving the labeling process;
- PMU and protective relay recordings with accurate timestamps are a better choice when determining and assigning event labels than SCADA information that has inherent limitations
- The physics-based model simulations of the power system can significantly enhance the ability to achieve efficient and reliable ML model training.
- To make a business case for the use of ML for event analysis it is crucial to compare the automated ML benefits over the manual approach to the analysis
- The ML approaches have a constraint in the developing power systems that may experience events never seen before, so keeping the history of new events for training purposes is essential

ACKNOWLEDGEMENT AND DISCLAIMER

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000913. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

BIBLIOGRAPHY

[1] Department of Energy, 10 CFR Part 1004, RIN 1901–AB44 "Critical Electric Infrastructure Information; New Administrative Procedures," Federal Register / Vol. 85, No. 51 / Monday, March 16, 2020 / Rules and Regulations.
[2] The Western Electricity Coordinating Council-WECC "Western Interconnection Data Sharing Agreement", 2016
[3] Electric Information Agency, "U.S. Electric System is Made up of Interconnections and Balancing Authorities,", July 2016
[4] IEEE Std C37.118.2TM-2011, "IEEE Standard for Synchrophasor Data Transfer for Power Systems," IEEE, 2011
[5] R. Baembitov, et. al. "Fast Extraction and Characterization of Fundamental Frequency Events from a Large PMU Dataset Using Big Data Analytics," HICSS-54 Conference, Hawaii, USA, January 2021.
[6] NERC, "Reliability Guideline: Forced Oscillation Monitoring & Mitigation," September 2017
[7] NERC, "Interconnection Oscillation Analysis: Reliability Assessment," July 2019
[8] A. Abdel Hai, et. al. "Transfer Learning for Event Detection from PMU Measurements with Scarce Labels" IEEE Access, Vol. 9, 127420 – 127432, 10 September 2021, DOI: 10.1109/ACCESS.2021.3111727
[9] M. Pavlovski, et. al. "Hierarchical Convolutional Neural Networks for Event Classification on PMU Measurements," IEEE Transactions on Instrumentation & Measurement. Vol. 70, 2021, DOI: 10.1109/TIM.2021.3115583
[10] L. Breiman, et al., Classification and regression trees. CRC press, 1984
[11] K. P. Murphy, "Logistic regression," Machine learning: a probabilistic perspective. Chapter 8, pp. 245–279, MIT press, 2012
[12] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," Nature, vol. 323, no. 6088, pp. 533–536, 1986