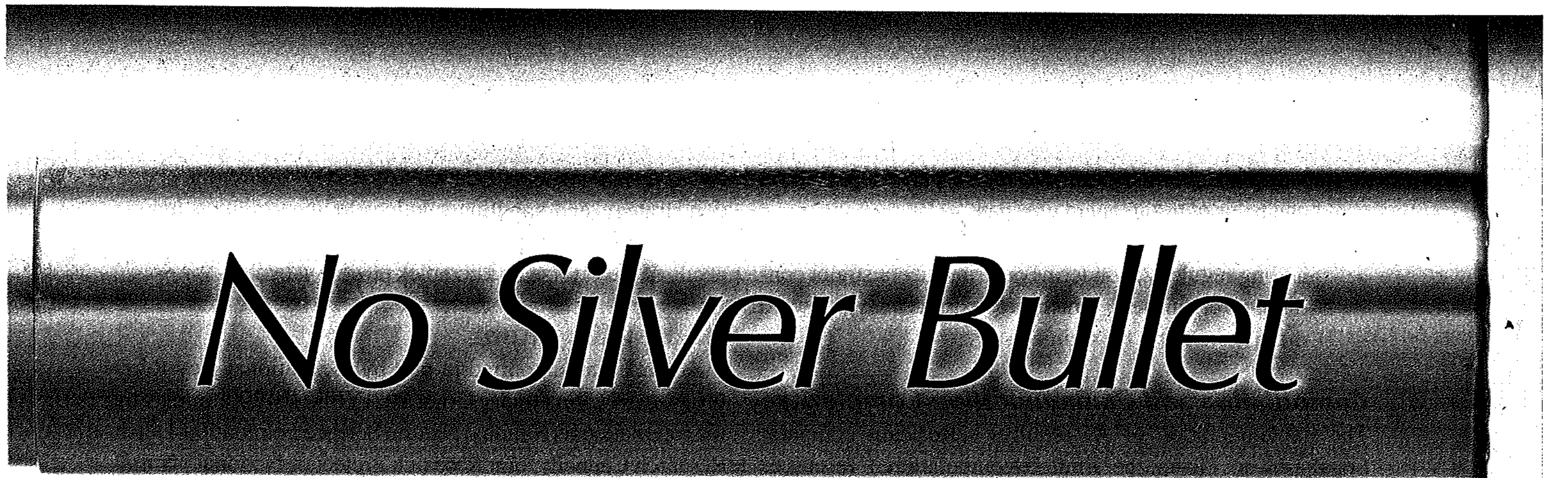

No Silver Bullet

TO FACILITATE REPAIR OF THE SYSTEM AFTER faults occur or implement appropriate control action to restore the power system's normal operation, the occurrence of faults is tracked through a variety of recording equipment, such as digital fault recorders (DFRs), protective relays, circuit breaker monitors, sequence-of-events recorders, fault locators, smart meters, power quality meters, and many other specialized devices; these are often located in substations or mounted on related power apparatus being monitored. A wealth of data is collected in real time, and additional descriptions of causes and impacts are recorded by maintenance crews as they inspect the sites where faults occurred. Such data are analyzed by utility staff to understand causes and impacts and develop mitigation measures. In general, this process takes time, and fault repair and restoration may be prolonged, preventing utility personnel from reducing or avoiding impacts, which, quite often, may be significant. In worst-case scenarios, the impact may be catastrophic, including loss of life.

This article focuses on the use of artificial intelligence (AI) to automate fault analysis, which quickly detects, classifies, and characterizes faults, resulting in an opportunity for timely actions to reduce fault impacts. The new opportunities of AI applications are made possible due to advances in data and computational sciences and new technologies for data collection, communication, and processing; above all, though, this possibility exists due to the capability of AI to perform causal analysis of fault-related data well beyond the cognitive capabilities of individuals or traditional analysis methodology. We describe how such techniques may be developed, what it takes to adjust current utility practice to implement them, and who may be the users of such solutions and related results.

Automated Fault Analysis in Transmission Systems

Issues With Existing Solutions

Traditional approaches to automated fault analysis in transmission systems focus on the use of data recorded in transmission substations. The typical data types and properties of substation recording equipment are shown in Table 1. Because such recording devices capture different data for different purposes, the users of such data are different groups in the utility company.

Close analysis of data types and properties from Table 1 differentiates the uses and users of data. Protection engineers use data obtained from triggered instruments, such as DFRs and digital protective relays, and obtain reports in the form of waveform and contact data samples. The state of the art is that such data are collected automatically whenever the faults occur and then analyzed automatically, allowing protection engineers to characterize faults quickly, initiate appropriate repair actions by informing repair crews where faults are located, and inform power system operators what types of faults occurred and whether or when the transmission lines that experienced faults may be restored. The disadvantage of such equipment is that the occurrence of faults triggers it and captures only a portion of the event data, starting with the trigger point and ending with fault clearance. The issue is that such data do not offer an extensive view of the disturbances describing the events that may be leading to a fault, including frequency excursions or disturbances caused by switching of power system equipment, including drooping of generators, shedding of load, or transmission line switching.

By Mladen Kezunovic, Rashid Baembitov, and Taif Mohamed

Digital Object Identifier 10.1109/MPE.2024.3415568
Date of current version: 12 November 2024

1540-7977/24©2024IEEE

Artificial Intelligence Is Not a Panacea, but It Works for Fault Analysis and Outage Management

Power system operators, on the other hand, get data from all substations using supervisory control and data acquisition (SCADA), but such data are somewhat limited by the recording features of remote terminal units (RTUs). RTUs report only root-mean-square (RMS) values and contact changes, which happen by exception when certain changes in the measured values exceed thresholds. A typical scan of SCADA systems happens only every 2–10 s, to help reduce the large amount of data being recorded. Hence, the events occurring between scan periods may not be captured. Besides, RMS values without the phase angle between the recoded three-phase waveforms of voltages and currents cannot be used to analyze fault type or location. That is why SCADA data are not used to analyze faults but only allow system operators to observe switching actions caused by relays.

Advantages of Using Phasor Measurement Unit Data

Phasor measurement unit (PMU) data are a relatively new addition to the recording capabilities in substations. The key features of PMU data are time synchronization and absolute time stamps.

Power system operators can get a much better resolution of an event recording since PMU data are typically reported 30–60 times a second and contain a detailed representation of the positive sequence phasors of current and voltages calculated from samples taken by PMUs. Since three-phase phasors and contacts may be captured by PMUs, protection engineers may use such data for fault analysis as well. Engineering groups may also use such data to analyze various power system phenomena not observed before, such as low-frequency oscillations caused by using renewables, and they can also use PMU data to verify models of power system components, such as generators and their controls. Because PMU data have such diverse and advanced properties, a business proposition in many companies has been developed around multiple uses of such high-fidelity data, which has led to their installation across the United States and worldwide at a rather intensive pace. It is estimated that over 4,000 PMUs may have been installed over the last 10 years in the U.S. system alone, and the number is increasing.

In the case of fault analysis, using PMU data helps system operators to quickly detect and classify faults without waiting for protection engineers to provide results of their analysis based on DFR and distribution point unit data. This is particularly important for the independent system operators that do not have access to DFR and digital protective relay data since the information is owned by transmission operators and not readily exchanged through the Open Access Same-Time Information System.

The uses and applications of PMU data do not come without major challenges, which eventually lead to the need to find a cost-effective way of dealing with such challenges to fully utilize the power of PMU data and justify the business case of investing further in this technology. Some major challenges include 1) an overwhelming amount of data coming from the PMU streaming capability, causing difficulties in processing large amounts of data by utility staff; 2) sparsity of PMU-equipped substations, causing difficulties in locating events and capturing their properties if the events are occurring in parts of the system far from the substations where PMUs are located; and 3) imperfections in data recording and communication systems, causing bad data and a lack of labels that describe how the PMU data patterns correspond to typical power system events, such as transients caused by faults, abnormal frequency deviation, or instability.

Such PMU data, coming from many substations, may have to be captured by multiple phasor data concentrators (PDCs) before it reaches the corporate PDC or super(regional) PDC, where various data storage capabilities are available and

table 1. Types and properties of substation-recorded data.

Source	Type of Data	Operating Mode	Reporting Frequency	Synchronization Among Devices
RTU	Status RMS value	Report	Every several seconds	No
PMU	Phasor	Report	Up to 60 times per second	Yes
ERD	RMS/DFT/samples	Upon request	Upon request	No

RTU: remote terminal unit; RMS: root mean square; PMU: phasor measurement unit; ERD: event reporting device; discreet Fourier transform.

various applications may reside. One simple observation is that the amount of streaming PMU data may be overwhelming for both humans and IT solutions currently processing such data. This, in turn, may create diminishing technology returns since the data may be rich in information, but timely knowledge extraction becomes difficult, if not impossible, limiting the full use of this technology. AI technology has become a niche solution since it enables automation of the analysis of large amounts of data very quickly, which is far more efficient than what humans or existing IT solutions may provide. At the same time, the knowledge extracted by AI solutions becomes readily available for all the PMU data users, particularly system operators, for their timely decision making. Further details are discussed in the upcoming sections.

AI Approach: From Data to Model

Machine learning (ML) techniques, as an AI approach, helped several fields leap ahead in terms of saving time and money to perform repetitive but complex tasks. A similar leap may be achieved in fault analysis if there is some guarantee that enough PMU data are available and that these data are dependable. Data availability is not simply achieved by collecting terabytes of PMU measurements but by ensuring that these measurements are diverse enough to accurately represent transmission system events. The power system mostly operates in normal conditions with some noise and subtle shifts in topology. Therefore, several hours or even days of PMU data may contain only normal operation, which is not enough when trying to train an ML algorithm to detect faults in the system. Detecting and classifying faults requires many instances of each type of fault targeted by the algorithm as well as a balanced number of normal operation instances to help the algorithm recognize the difference.

When data are collected at such a large scale, the quality of the collected data may be challenging to maintain. A PMU, like any other device, may encounter an error or go offline for a multitude of reasons. Because fault detection algorithms are looking to flag anomalies in data, bad data may be mistaken for a critical fault event. Missing and unreasonable data points, sometimes intervals, are the most difficult to deal with. In some big data applications, the issue of missing data may be mitigated by traditional methods, such as interpolation or advanced techniques, such as variational autoencoders or generative adversarial networks. However, when the main goal is to detect abnormal behavior, trying to fill in any of the missing data intervals with reduced accuracy may lead to misclassifying an event as a normal operation. That is the point where the role of domain expert judgment becomes most apparent. For example, a PMU with large amounts of missing data or unreasonable data from recordings in an actual system for several years may be disregarded from the process. In addition, the features chosen to train an ML algorithm may be modified to consider this phenomenon. Part of evaluating PMU data quality is evaluating the quality of the accompanying event logs. Event logs are descriptions of power system events, one category of which is line faults. Event logs contain labels that define when an event occurred through a time stamp, what type of event it is, and any further classification or characterization of the event. Labels are extremely important for developing a supervised ML model for fault detection and classification.

Labels exhibit several problems that include inaccurate time stamps, wrong event type descriptions, and wrong fault type classification. While visual inspection of large amounts of data may be a tedious process, it may be required to create accurate labels, which are a precursor for effective AI model training.

One way to produce more event labels and ensure their accuracy is through simulation. In the case of line faults, some faults are more commonly found in field-recorded PMU data than others. Three-phase faults, for instance, are rare in comparison to single-phase-to-ground faults. Nonetheless, simulating more three-phase faults using a system like the one portrayed in Figure 1 achieves a balance in data, giving the classifier model enough training information with accurate labels to distinguish among different fault types.

Features and Model Implementation

Feature engineering is one of the most important steps in designing a successful ML model for fault detection and classification. It is the step in which one decides how to present the information to the ML algorithm.

Despite the ability of AI tools to detect patterns and point out irregularities, domain experts know why and where to look for these irregularities; this is the advantage of feature engineering done by domain experts. In the case of line faults, dips in voltage are the main indicator that a fault has occurred. Voltage dips signify that a relay has engaged the circuit breakers on the line and are usually followed by an autoreclosing sequence to automatically clear the fault and return to normal operation. Consequently, a feature capturing such voltage dips would be the best candidate for training an ML algorithm. To provide a clearer picture of the expected data, Figure 2 shows several field-recorded PMU measurements of single-phase voltage over the intervals of 1 min [Figure 2(a)] and 2 s [Figure 2(b)]. To the untrained eye, it might seem like there are numerous voltage dips. However, what is displayed are the normal conditions under which the power system operates, where there is a lot of noise. There are two important takeaways from this figure. The first is the importance of the use of field-recorded data in the training process of ML models since this behavior is not captured by simulations. The second is the importance of choosing the time window for extracting features from the data. A time window of 1 min may be too long and may contain too many perturbations for a fault detection algorithm. A 2-s time window, on the other hand, allows for precise detection of the time at which the fault occurred. Realistically, faults cause transients that occur at the scale of milliseconds. Even though using such a short time window is possible, it is not practical considering the amount of data available that span months and years. The most advantageous time window, such as the 2 s selected in this application, would be a compromise.

After extracting features from the PMU data, both simulated and field-recorded, the model is brought together by selecting the ML algorithm and using the extracted features as input. Using supervised learning techniques, the model is trained to associate certain features with the labels it receives for their respective time windows. Accordingly, the process of label improvement by domain experts is also crucial to the model's resulting accuracy.

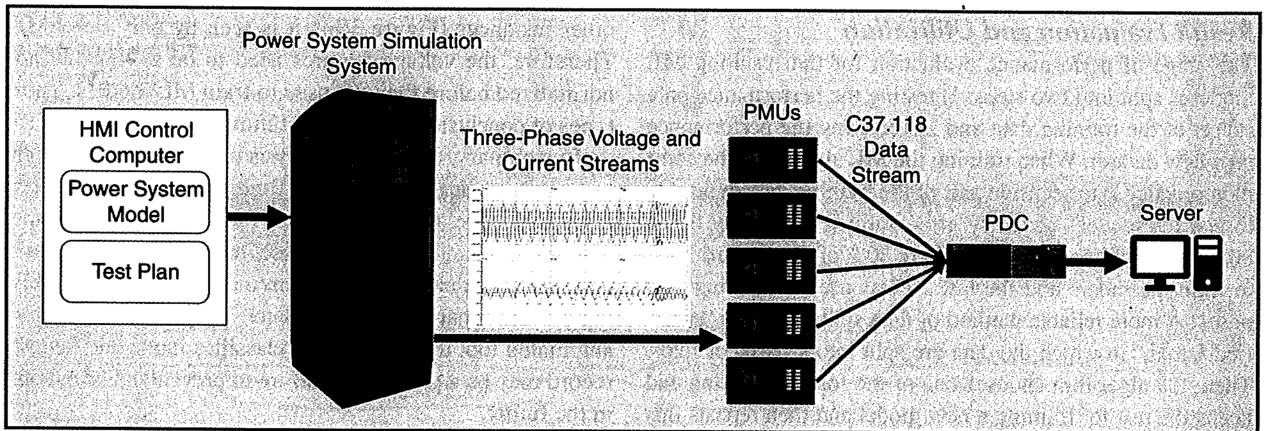

figure 1. The process of creating fault events through simulation. HMI: human–machine interface.

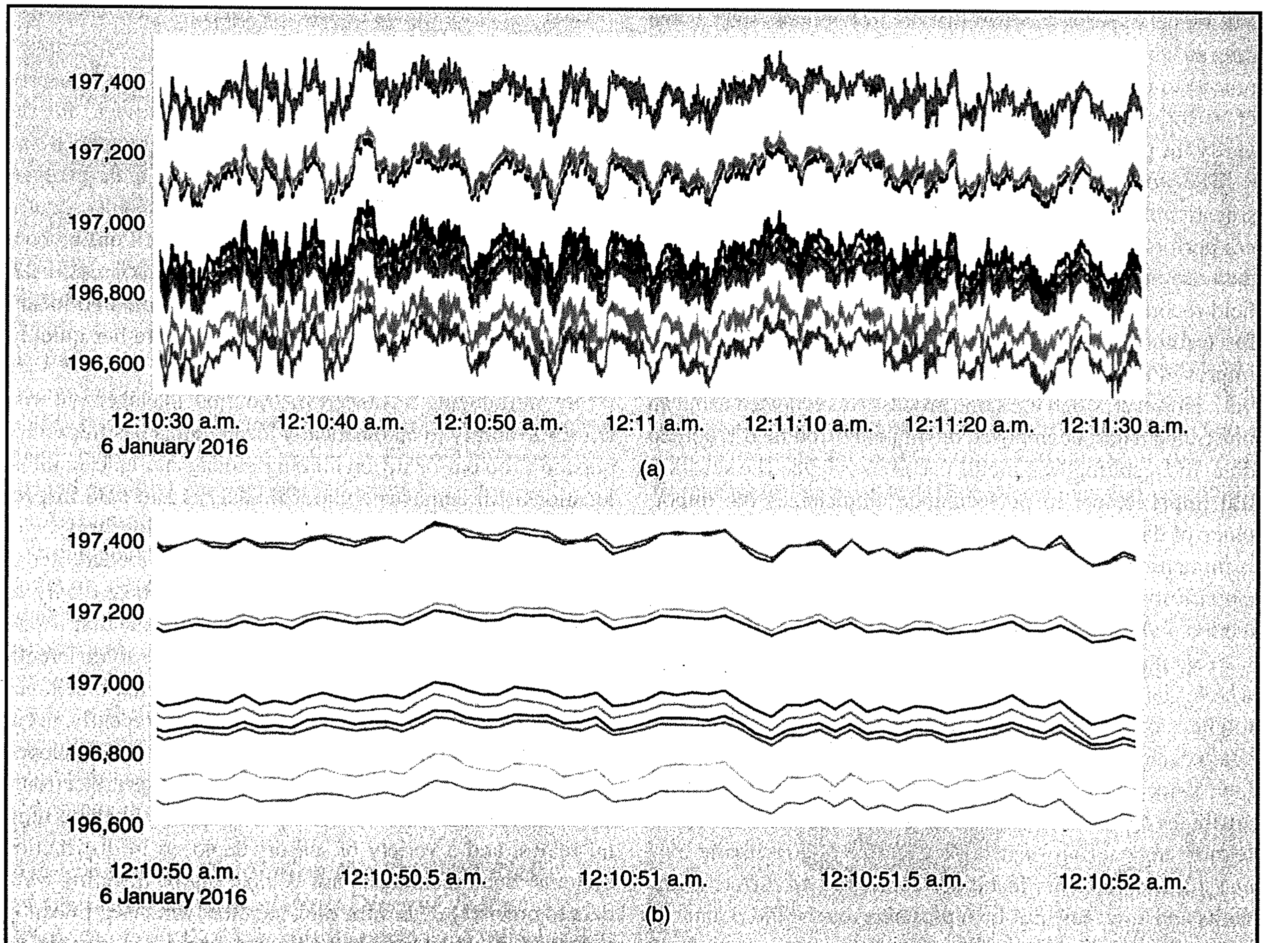

figure 2. PMU ambient data for (a) 1 min and (b) 2 s.

The scarcity of relevant examples in data can hinder the potential of an ML algorithm; this translates into the need for well-documented structured data.

Result Evaluation and Utilization

The stage of performance evaluation for the resulting ML model is split into two steps: 1) testing the performance on a subset of the training data and 2) verifying the performance on a new dataset. When training the ML algorithm, the common practice is to set aside part of the training data to be able to test the model's accuracy after training. This may be an 80/20 data split, for example, where 80% of the data are used to train the model and the last 20% is used for testing purposes. A more reliable method of data splitting is the "stratified k-fold," in which the data are split into k parts, or folds. Then, the algorithm chooses one of the folds for testing and keeps the rest for training a new model and then repeats this process k times. This process ensures that the model's performance is cross validated over several train/test datasets independently. Training and testing may also include fine-tuning ML model hyperparameters that control options, such as the rate at which a model learns and its complexity (strongly related to the model's overfitting). When an acceptable level of performance is reached by the model, the next step is to verify the performance using a new set of field-recorded data.

The process of evaluating several fault detection and classification models has consistently shown that these classifiers perform better when both simulated and field-recorded data are used for training. In the best-case scenario, using field-recorded data only, when all the time windows used for training were inspected visually and the labels were improved accordingly, the resulting weighted F_1 score was 0.87. However, when the same process was repeated using an integrated dataset composed of simulated and field-recorded data, the resulting weighted F_1 score was 0.98. The substantial improvement in performance emphasizes the importance of data balance when developing these models. There is immense potential for fully automated fault detection and classification in the field with the right practices of data collection and categorization.

Yet, the best model did not achieve 100% accuracy, which sparks interest in understanding where these models might make mistakes. The domain expert's insight is once again essential at this point. Inspecting the results of the line fault detection and classification algorithm emphasized the attributes of line faults that need further attention during the feature engineering stage. One example of a recurring misclassification appears in Figure 3, where two different line faults and how they may be represented by the PMU data are illustrated. While line faults are localized events, in some instances [Figure 3(a)], a fault is detected by all PMUs; in other instances [Figure 3(b)], it is seen by only one PMU. Therefore, the voltage features need to be aggregated and normalized before they are used to train ML models. These types of conclusions help guide future studies.

In summary, automated analysis of faults using AI techniques dramatically reduces the time and effort that must be spent to analyze faults. Even when faults do not need immediate maintenance attention, a record of reoccurring faults on a line or in an area may be an indication of an underlying problem that may escalate later. As a result, having an automated tool that detects and classifies faults and keeps a record may be a proactive measure to prevent major outages in the future.

Automated Prediction of Outage State of Risk in Distribution Systems

Issues With Existing Solutions

Outages in distribution systems occur relatively frequently, but the impacts may not be noticed by consumers for the following reasons: 1) the outages may be mitigated by an autoreclosing action of circuit breakers aimed at de-energizing the system briefly to eliminate temporary faults, 2) utility personnel take quick action using advanced online control means to restore power from alternative sources, or 3) utility staff locate the fault accurately by automated means and send crews to repair and restore system rather quickly, in the case of permanent faults.

By introducing advanced distribution management systems, the ability to automatically locate faults has improved, but the problem of differentiating among multiple possible locations still demands additional analysis and field inspection. It is especially challenging in distribution grids with radial topology, due to the high density of feeders and a lack of distributed sensors. Measurements taken only at the feeder root (substation) may render a result that shows a fault at multiple locations, due to an inability to differentiate which feeder and/or lateral created the electrical distance (impedance) between the root and the fault. Recently, accurate automated fault location techniques have been developed based on data collected from different intelligent electronic devices, such as smart meters, power quality meters, digital relays, and a variety of sensors deployed in distribution systems. Such techniques use voltage sag or traveling-wave data to precisely determine fault location, which significantly reduces mitigation time by utility personnel. Despite all the technological advances, the current practice in utilities is to

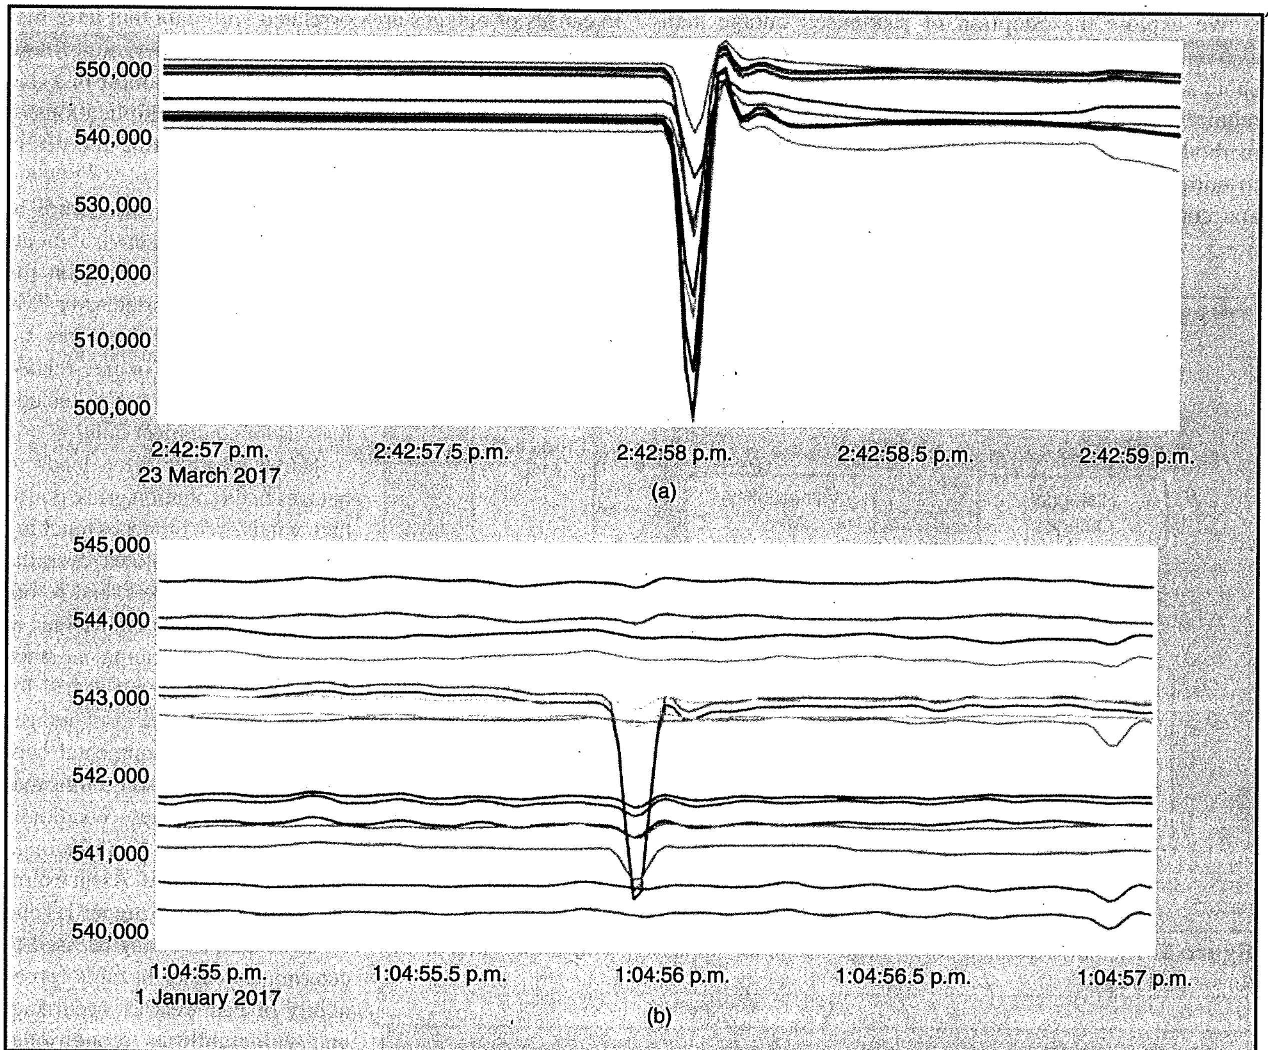

figure 3. A fault seen (a) by all PMUs and (b) by only one.

perform reactive measures that are deployed only after the outage has occurred. This emphasizes two drivers of the reliability issue in distribution system operation: utilities focused on revenue and assets and customers focused on power quality and needs, as summarized in Figure 4.

As Figure 4 demonstrates, the existing practice of reactive outage mitigation and management separates the utility focus on system restoration and asset management as major goals driving system recovery; this leaves the issue of power quality and customer needs to customers to deal with during outages. Such practice does not take advantage of predictive methods that, if feasible, would allow utilities to not only focus on restoring the system after outages occur but also help consumers by giving them options on how to mitigate the outage impacts until the service is restored.

Advantages of Predicting the State of Risk of Outages by Correlating Diverse Data

The state of risk (SoR) for distribution outages reflects how likely an outage may be in a certain location, at a predefined time interval. Outage SoR prediction leads to a proactive approach, as opposed to a reactive approach; rather than acting only after the outage has occurred, utilities and consumers may take mitigating actions beforehand, which leads to managing the adverse effects of outages and potentially reducing losses.

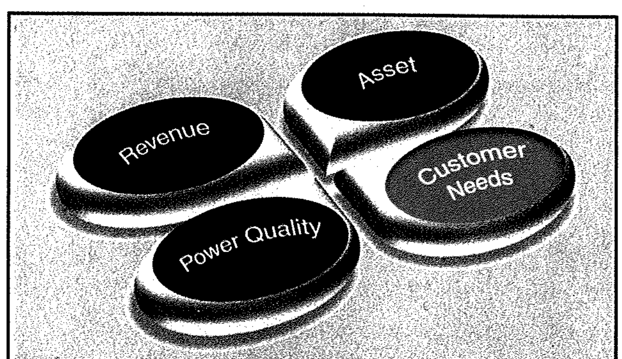

figure 4. Different assessments of outages.

We explore the adoption of risk-based outage management, as it enables different customers to be notified about a potential outage in advance. Such an approach may require two-step action: 1) predicting outages using AI methods and 2) developing management strategies on how to mitigate fault impacts by reducing the risk on utilities and consumers. This new paradigm is illustrated in Figure 5, where weather and other environmental data relating to causes of outages are correlated with data that have historically been related to outages and then merged through ML/AI techniques to create risk maps (an example of a risk map is presented in Figure 6), which are multiplications of hazard maps and vulnerability maps reflecting the direct causes of outages.

To predict outage occurrences in the grid, one needs to associate them with causes. In recent years, environment-related outages have been on the rise. As widely reported, over 75% of distribution system outages are caused by weather or tree-related causes which, in turn, also include reasons for equipment failures.

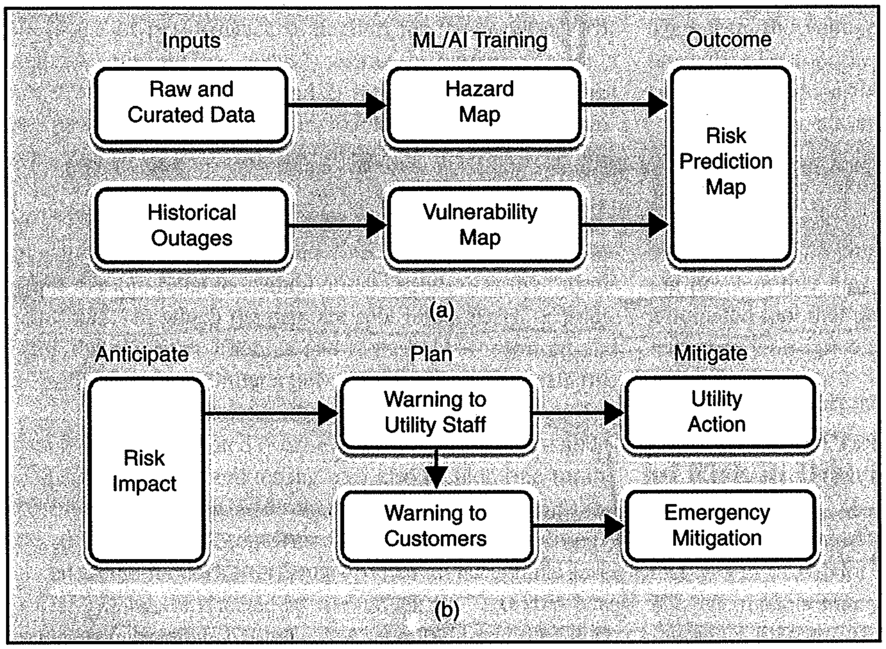

figure 5. The new two-step SoR outage prediction approach. The (a) first step and (b) second step.

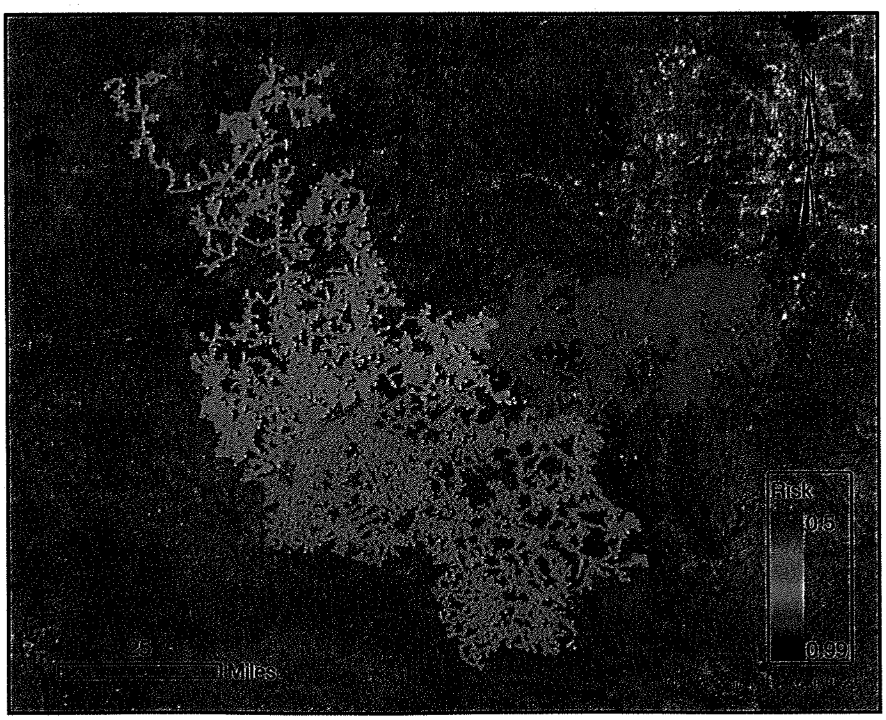

figure 6. The outage risk prediction represented in a geographic information system.

Harsh weather often leads to outages in distribution grids, due to high winds, tree branches touching the wires, debris blown over the lines, lightning strikes, and so on. Also, harsh weather contributes to wear and tear on equipment since most of the equipment is still located outdoors. Since most outages are related to environmental conditions, the use of some information that reflects these conditions is needed, and then a correlation to outages is performed. As an example, Table 2 lists the properties and types of data that may be used to determine weather hazards (probability of bad weather occurring) and vulnerabilities (conditional probability of harsh weather causing faults) that, when multiplied, create the SoR estimate. The SoR reflects the probability of an outage. Further details of how such data may be correlated to calculate hazards and vulnerabilities using AI to predict the SoR of outages are presented next.

AI Approach: From Data to Model

Weather forecasts and historical weather records are the core data sources for outage prediction models. With the advance of new-generation weather models, atmospheric conditions can be predicted in a very detailed way. Several relevant weather parameters become readily available for use. These include, but are not limited to, wind

table 2. Types and properties of weather and other environmental data that may be used for SoR prediction.

			Velocity	Volume	
	Source	**Data Type**	**Temporal Resolution**	**Spatial Resolution**	**Measurements**
Variety	Automated Surface Observing System	Land based	1 min	900 stations	Air temperature; dew point; relative humidity; wind direction, speed, and gusts; altimeter; sea level pressure; precipitation; visibility
	Level-2 next-generation weather radar	Radar data	5 min	160 high-resolution Doppler radar sites	Precipitation and atmospheric movement
	NOAA satellite database	Satellite data	Hourly, daily, monthly	4 km	Cloud coverage, hydrological observations (precipitation, cloud, liquid water, total precipitable water, snow cover), pollution monitoring
	Vaisala National Lightning Detection Network	Lightning data	Instantaneous	Median location accuracy: 200–500 m	Date and time, latitude and longitude, peak amplitude, polarity, type of event (cloud or cloud to ground)
	National Digital Forecast Database	Weather forecast data	3 h	5 km	Wind speed, direction, and gusts; relative humidity; convective hazard outlook; tornado probability; thunderstorm probability
	Texas Parks and Wildlife Department	Texas ecological mapping systems data	Static	10 m	Distribution of different tree species
	Texas Natural Resources Information System	NAP	Year	50 cm — 1 m	High Resolution Imagery
	NASA	3D global vegetation map	Static	1 km	Canopy height data
	National Cooperative Soil Survey	Gridded Soil Survey Geographic Database	Static	10 m	Soil type
	Utility	Historical outage data	Instantaneous	Feeder section	Location, start and end time and date, number of customers affected, cause code
		Tree trimming data	Day	Feeder	Feeder location, date, trimming period, number of customers affected, cost of trimming
		Network geographic information system data	Static	Infinity (shapefile)	Pole location, material/class, and height; feeder location, conductor size, count, and material; nominal voltage
		Historical maintenance data	Day	Tower location	Start and end date and time, location, type (maintenance, replacement), cost, number of customers affected
		Insulator asset data	Static	Infinity (shapefile)	Surge impedances of towers and ground wires, footing resistance, component basic insulation level
		In-field measurements	Instantaneous	Tower location	Leakage current magnitude, flashover voltage, electric field distribution, corona discharge detection, infrared reflection thermography, visual inspection reports

It is harder to predict accurate SoRs for a smaller area of focus on a long-term horizon than it is for a large territory a couple of hours ahead.

direction and wind speed, precipitation levels, temperatures at several altitudes, atmospheric pressure, relative humidity, cloud cover, and the Lightning Potential Index, among others.

The relevant datasets are not limited to weather. Several other factors play a role in distribution outages. Vegetation data provide insights into where trees are located and how close they are to power lines. The types of trees can help estimate their growth rates when combined with precipitation and temperature data. The underlying soils influence vegetation growth, and relevant data on soil types can also be obtained and used.

Direct lightning strikes to phase wires cause high voltage that translates to insulation flashovers. Also, indirect lightning strikes near power lines induce a high voltage in the lines, causing outages. For these reasons, lightning detection network data have become useful in outage management. Several commercial vendors provide data for lightning strike locations. The data include several parameters: time of a lightning strike, type of lightning strike, maximum current, time rise, and so on. If properly processed, these variables can be utilized in ML models for outage prediction.

Once all the available relevant data are gathered, the next step is to put them all together, considering the difference in the properties, as described in Table 2. The process is referred to as *spatiotemporal correlation of data* and is meant to identify the environmental conditions around the network at the time when an outage has happened in the past (i.e., capturing what the weather and other conditions were like when an outage occurred).

These occurrences (examples) are referred to as *training datasets*. The underlying idea is that as historical outage examples are gathered, information about what has led to an outage is accumulated. For instance, an outage can tend to occur when the wind speed over an area with power lines is higher than 18 m/s. However, to give the model information about normal operating conditions, several points in time when no outage has occurred are also selected. This information complements the training dataset. Ideally, one would need a balance between outage examples and examples of normal operations. In this case, the data model would have a better differentiating ability; that is, it would classify a point in time as an outage or no outage more accurately.

The next step is to wrangle the data into a suitable form for ML algorithms. To be used by the ML model, the data go through cleansing, scaling, and transformation in a sequential order. This way, the data are wrangled to obtain a processed, structured, and clean training dataset.

Features and Model Implementation

Once the training dataset is prepared, it is fed to the ML algorithm. The algorithm then "learns" the intrinsic dependencies between outage occurrences and the relevant data. The benefit of ML is that, given sufficient examples, it can infer the knowledge hidden in the data. The benefit of ML over classic deterministic approaches is that ML algorithms allow for understanding phenomena by using data that reflect what actually happened in systems. This augments the analysis by utilizing predetermined scenarios, which is the approach taken in the analysis that is based on classical physics-based models. If enough data are gathered on the phenomena (in our case, power outages) and an underlying correlation exists (environmental impacts lead to outages), then the ML would be able to "dig out," or mine, the knowledge without human intervention. At the same time, the benefit also creates a limitation: the quality of the ML model can be only as good as the data it was trained on. Subsequently, the scarcity of relevant examples in data can hinder the potential of an ML algorithm; this translates into the need for well-documented structured data. We advocate for adopting a data-oriented culture in every utility.

Critical features of an outage SoR prediction are its spatial and temporal applications. The outage SoR predictive model can be adopted for a variety of temporal horizons and spatial resolutions. For immediate and short-term actions, as an example, a model may be trained to predict outages in the next 24–48 h by utilizing a state-of-the-art weather forecast. For long-term planning, a model may be trained using prevailing climate conditions to evaluate the system's susceptibility to upcoming seasonal weather changes. On a spatial dimension, the model can output SoR prediction estimates for the entire service territory or focus on a specific part, such as a single feeder or feeders connected to the same substation. The selection of a spatial and temporal resolution comes with considerations for the accuracy of predictions. It is harder to predict accurate SoRs for a smaller area of focus on a long-term horizon than it is for a large territory a couple of hours ahead. The model's performance is proportional to the spatial resolution and inversely proportional to the temporal horizon.

The balance between spatiotemporal resolution (i.e., how far in time one needs the predictions and how spatially granular they need to be) drives the accuracy of the model for each application. Spatiotemporal resolution is also dependent on the mitigation actions that each stakeholder is willing and able to take during a certain period and on a given

The models are not perfect and make mistakes, which is expected, given the probabilistic nature of outage occurrence and data imprecision or uncertainty.

spatial scale. Leaving too short a time to react diminishes the value of SoR outages, while having too broad an area of focus might limit the effectiveness of mitigation actions. The optimal spatiotemporal resolution is achieved through several iterations of adjustments in a period when the utility and customers provide their feedback so that the needs and limitations are incorporated into the solution.

Result Evaluation and Utilization

Once the model is trained and optimal temporal and spatial resolutions are agreed upon, one needs to evaluate the potential performance of the SoR prediction model. The standard way of testing an ML application is to see how it performs on data that it has never been given before. In such a manner, it is ensured that the evaluation of the performance is carried out objectively. Testing ML algorithms on data that were used during the training phase produces overoptimistic results since the algorithm is already optimized for such data.

To quantify the performance of models, researchers use a different set of metrics. Each metric measures a specific ability of a model and is not as useful when used in isolation from other metrics. For example, the precision metric shows a measure of events that are correctly classified as having an outage out of all the events with outages. If that had been a single metric used for evaluation, the model could have classified all the events as an outage and achieved 100% precision! Hence, the model would be useless. That is why the metrics are used in combination with one another, as described next. Given a balanced dataset, the precision during the testing phase can reach up to 78%, with a corresponding recall of around 67%.

One can think of a data model as a decision-making machine: it answers the question of whether a given period has been an outage. The models are not perfect and make mistakes, which is expected, given the probabilistic nature of outage occurrence and data imprecision or uncertainty. The mistakes, or errors, of the models can be of two types: 1) a type 1 error, or a false positive, appears when an event is classified as an outage when there was no outage, while 2) a type 2 error, or a false negative, happens when an outage is classified as normal operation. The importance of differentiating between the two is in the unequal costs of errors for different types of customers and utilities. For instance, it might appear excessive to dispatch portable generation into a residential area when the risk of an outage is high. However, the same elevated risk levels, if threatening the operations of a hospital, may require ensuring a secure power supply by any means.

To coordinate mitigation actions on the utilities' side, we created a tool allowing personnel to have access to the prediction of SoR levels. The tool creates a visual representation of outage SoR levels overlapping a map of the distribution feeders. Such geographic representations of SoR predictions are referred to as *risk maps*. An example of a risk map is provided in Figure 6. The map conveys the conditional probability of outage occurrence in a given region and time, subject to a probability of given hazardous environmental conditions affecting the grid during that time. This may also be viewed as a multiplication of the two probabilities. The risk analysis can be much more complex by introducing the probability of other causal impacts, such as fire, animal intrusion, or human negligence, but such complexity is beyond the scope of our study at this time. Since we are dealing with probabilities that reflect uncertainty about whether an event will occur or not, the extremes of zero and one are reflective of the potential that the event never occurs or always occurs.

Given a set of risk maps for different time horizons, utility staff can effectively select and execute necessary mitigation measures. The measures can be differentiated for different parts of the system, based on the potential losses from outage occurrences. At the same time, risk maps allow utility personnel to understand the overall risk levels in the distribution grid. Such information can help plan the number of repair crews for the day ahead. For example, if the risk levels stay elevated after 5 p.m., the staff coordinator notifies the crews on call that they will likely be called for repairs that evening.

The predicted SoRs are also useful for alerting customers about potential outages in their area. Given an advance notification, customers may adjust their plans and prepare for possible outages according to their exposure to outage impacts. Big industrial customers may suffer heavy losses from unproduced goods or restarts of power-sensitive production lines. Small businesses lose customers if an outage happens during business hours. Incurring reputational damages may affect business in the long run. For example, food spoilage from disrupted refrigeration exacerbates the problem.

In summary, outages significantly disrupt everyday activities, leisure time, planned meetings, and overall quality of life, but SoR prediction allows utilities and consumers to plan mitigation measures ahead of time and proactively manage the risk; this changes the level of the outage impacts. The rise of remote work postpandemic intensifies

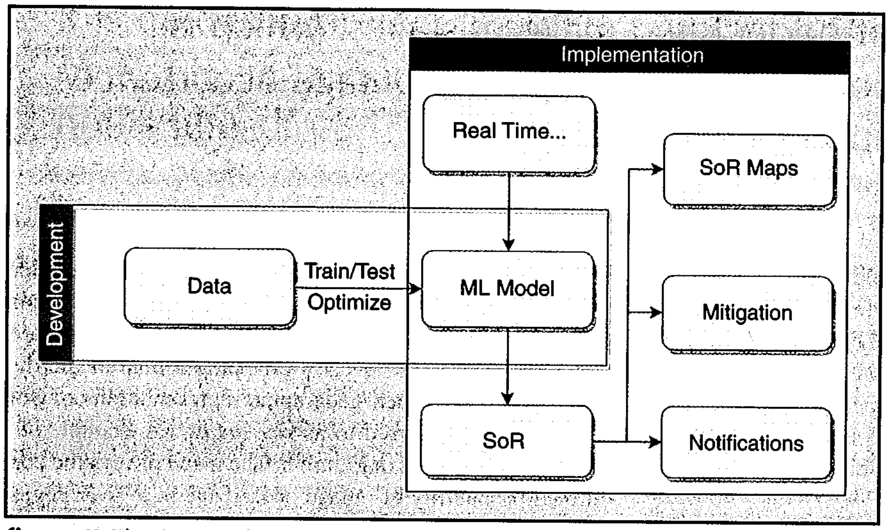

figure 7. The SoR prediction application.

the direct impact on an individual's ability to perform work duties, potentially hindering career growth and leading to unemployment. Vulnerable populations, such as those dependent on power for medical appliances and the elderly or handicapped, are particularly affected by outages, which can profoundly impact their lives. All such negative impacts can be significantly reduced because of the ML use for SoR outage prediction.

Research findings highlight that a simple 1-h advance notification to customers regarding potential outages can result in a substantial reduction in outage costs in some cases. Notifying customers about potential outages transforms an unforeseen and forced event into a planned occurrence with an expected probability.

Implementation and benefits of the proposed risk-based framework can be summarized in a simple diagram, given in Figure 7. Relevant data sources are correlated together to produce training and testing datasets for ML algorithms. The pretrained models are then utilized on real-time data input to produce outage prediction in the form of SoR maps. The SoR maps are then utilized to devise optimized mitigation actions to reduce the impacts of outages.

Conclusions

The use cases discussed in this article illustrate how AI may be used to handle the analysis of faults in transmission systems and the management of outages in the distribution system. In both cases, several advantages are recognized:

- ✔ AI allows for the development of automated solutions that correlate large amounts of data in a short period of time, a performance unmatched by humans.
- ✔ The AI approach enables operators, protection engineers, and asset and outage management crews to develop mitigation actions ahead of time, enabling a reduction of impacts.
- ✔ The utilization of data for AI applications requires corporate-wide attention due to the variety of issues related to data preparation and storage, which go well beyond current utility practice.
- ✔ The development of AI applications may require different skill sets, from data analytics to IT and communications, but it always requires extensive domain expertise.

Building a business case for the deployment of AI must include carefully weighing the obvious advantages of AI use against the cost of implementing such solutions as well as training personnel for utilizing the results.

For Further Reading

R. Baembitov and M. Kezunovic, "State of risk prediction for management and mitigation of vegetation and weather caused outages in distribution networks," *IEEE Access*, vol. 11, pp. 113,864–113,875, 2023, doi: 10.1109/ACCESS.2023.3324609.

T. Dokic and M. Kezunovic, "Predictive risk management for dynamic tree trimming scheduling for distribution networks," *IEEE Trans. Smart Grid*, vol. 10, no. 5, pp. 4776–4785, Sep. 2018, doi: 10.1109/TSG.2018.2868457.

S. Raschka, "Model evaluation, model selection, and algorithm selection in machine learning," 2020, *arXiv:1811.12808*.

Y. Cheng et al., "Real-world subsynchronous oscillation events in power grids with high penetrations of inverter-based resources," *IEEE Trans. Power Syst.*, vol. 38, no. 1, pp. 316–330, Jan. 2023, doi: 10.1109/TPWRS.2022.3161418.

F. Yang, D. Cerrai, and E. N. Anagnostou, "The effect of lead-time weather forecast uncertainty on outage prediction modeling," *Forecasting*, vol. 3, no. 3, pp. 501–516, 2021, doi: 10.3390/forecast3030031.

R. Tervo, J. Karjalainen, and A. Jung, "Short-term prediction of electricity outages caused by convective storms," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 11, pp. 8618–8626, Nov. 2019, doi: 10.1109/TGRS.2019.2921809.

Biographies

Mladen Kezunovic is with Texas A&M University, College Station, TX 77843 USA.

Rashid Baembitov is with Texas A&M University, College Station, TX 77843 USA.

Taif Mohamed is with Texas A&M University, College Station, TX 77843 USA.