Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 59-9 (2025) 181–186

Interpretability of the ML-Based State of Risk Predictions for the Electric Grid Forced Outages

R. Baembitov[*], A. Karmacharya[*], M. Kezunovic[*], D. Saranovic[**], Z. Obradovic[**]

[*]*Texas A&M University, College Station, TX 77843-3128*
USA (949-899-7214, bae_rashid@tamu.edu; abinash_k@tamu.edu; m-kezunovic@tamu.edu)
[**]*Temple University, Philadelphia, PA 19122, USA*
(daniel.saranovic@temple.edu; zoran.obradovic@temple.edu)

Abstract: The electric grid forced outages (further referred to as "outages") are increasingly common due to the growing complexity of the electric grid, wear and tear of the components, and the impact of inclement weather. The recent development of Machine Learning (ML) approaches for predicting the outage State of Risk (SoR) allows the time needed to implement mitigation measures that can reduce the risk and, consequently, impacts associated with the outages. This paper establishes a baseline for the outage SoR prediction using data from historical outage logs and relevant weather parameters associated with an actual power system in Texas, USA. The baseline is used to demonstrate how various design requirements for the outage SoR prediction models can impact the interpretability of the results. The paper's findings emphasize the implications of different interpretability outcomes that affect the effectiveness of the mitigation options.

Keywords: Machine Learning, Artificial Intelligence, Risk, Outage, Interpretability

1. INTRODUCTION

Since most of the outages are caused by inclement weather or equipment wear and tear, it was recognized a decade ago that using ML approaches on big data associated with the causes of outages may enable the prediction of outage State of Risk (SoR) (Kezunovic et al., 2013). Subsequently, many studies have explored different data and algorithms used to develop the outage SoR prediction. Our contribution reveals how the interpretability of the results may affect the mitigation decision-making of different end-users within the utilities and their customer base.

1.1 State of art in outage SoR prediction

Many machine learning algorithms such as logistic regression (LR), random forest (RF), gradient and adaptive boosting (GB, AB), and multi-layer feed-forward networks (MFN) were utilized for outage prediction using the historical weather and outage data, leading to the conclusion that the model's performance depended upon the area size (Garland et al., 2023). A five-step framework was developed encompassing failure probability prediction using LR to enhance system resilience for overhead lines due to lightning events (Mujjuni et al., 2023). However, the study only considered a single 400kV transmission line rather than a system network. Deep learning techniques were used to predict the outage probability in a census tract based on weather, infrastructure, and socio-economic details (Wang et al., 2024). A comparison of autoregressive statistical approaches was performed against time series-based deep learning models like Long Short Term Memory (LSTM) and 1D – Convolutional Neural Network (CNN) for predicting customer outages over varying lead times (Udeh et al., 2022). Different machine learning algorithms like RF, neural network (NN), support vector machine (SVM), k-nearest neighbors, and decision tree (DT) were used to predict and classify outage causes (Kor et al., 2020). These references consider very large-scale networks over census tracts or counties in contrast to our entire distribution scale network, over which the outage is much more difficult to predict accurately.

Graph Neural Network (GNN) has been utilized to predict outage occurrence based on the data obtained from weather stations and their relative location in the area of interest (Owerko et al., 2018). Survival model-based outage risk prediction has also been proposed by combining various resilience-based metrics with operation attributes and vegetation indicators (Jain et al., 2021). Fragility curves for weather-related outages were obtained through catastrophic risk modeling based on the wind speed values and outage data at varying spatial resolutions (Dunn et al., 2018).

1.2 Our prior work

Early on we explored Big Data applications in outage management (Kezunovic et al., 2013), and then applicability of ML algorithms for predicting outage SoR by incorporating graph embeddings (Baembitov et al., 2021). We then studied the influence of wind modeling (Baembitov et al., 2023) and lightning (Baembitov et al., 2025) on outage SoR predictions, and a sensitivity analysis of various ML algorithms (Baembitov et al., 2024). We also explored other aspects of outage SoR prediction: a). mitigation strategy for consumers once notifies of the imminent outages (Baembitov and Kezunovic, 2023), b) prosumer strategy optimization utilizing SoR levels (Khoshjahan et al., 2023), and c) tree-trimming plan to reduce the overall risk in the system (Dokic and Kezunovic, 2019).

In this paper, we are focusing on design parameters and results interpretability and emphasizing the importance of the interpretability of the model outputs for decision-making.es).

10.1016/j.ifacol.2025.08.133

1.3 Paper contribution

We could not identify references that demonstrate the variability in the ML result interpretability by the end-user. While some references offer a sensitivity study of the resulting ML model metrics, the uses of the results for decision-making regarding risk mitigation are not sufficiently explored. In this context, our contributions are:

- An analysis of the SoR prediction implementation options that may affect the results performance metrics.
- Development of the baseline for algorithm comparison along the spatiotemporal variability.
- Classification of the user decision-making needs and requirements as they interpret the outage SoR results.

2. DESIGN REQUIREMENTS AND IMPLEMENTATION OPTIONS FOR THE OUTAGE SoR PREDICTION

2.1 Typical data requirements and options

Since the causes of outages vary by various areas in the service territory, an optimal selection of data should represent outage causes for given circumstances of the grid disposition and environmental/weather exposures.

Several outage causes reported in the literature (Fig. 1) correlate to different data sources (Table 1). The key to merging such data for ML model development is correlating the data in time and space while accounting for data uncertainty, errors, and refresh rate.

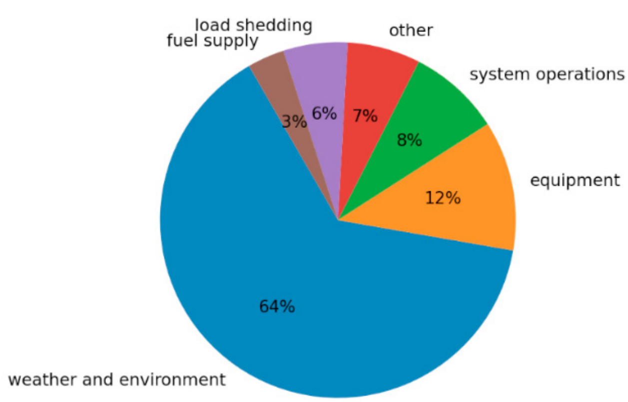

Fig 1. Outage causes in the U.S. (DOE, 2002-2021).

Table 1. Data types

Data Type	Description
Utility sources	
Feeder locations	Location of overhead lines in the network
Historical outage logs	A list of outages that occurred in the network in the past
Substation locations	Coordinates of the network substations
Historical tree trimming logs	Where, when, and how much trimming is performed
Historical maintenance logs	Where, when, and which equipment was repaired/replaced/maintained
Costs of outages	Outage monetary impact
Customer types and locations	Aggregated and/or anonymized customer groups
Public sources	
Historic weather condition logs	Meteorological parameters with timestamps and locations
Weather forecasts	Data reflecting future weather conditions
Vegetation types	Vegetation properties (foliage, canopy, height, crown)
Leaf Area Index (LAI)	Amount of leaf material in a plant canopy
Paid sources	
Light Detection and Ranging (LiDAR) scans	3D images of the terrain of interest
Historic lightning logs	Location, time, and type of lightning strikes in the area

2.2 Data/model requirements and typical algorithm choices

To illustrate how the data may be utilized, it appears convenient to represent data in one of the visualization software tools, such as ArcGIS, allowing a layered view of data superimposed on the grid topology (Fig. 2). The map was acquired using a Random Forest classifier; more information about models and their relative performance can be found in (Baembitov et al., 2024). The area with outage prediction is marked in red, the area where the outage is unlikely is marked in blue, and just for comparison purposes, the location where the outage occurred is marked in green (ground truth).

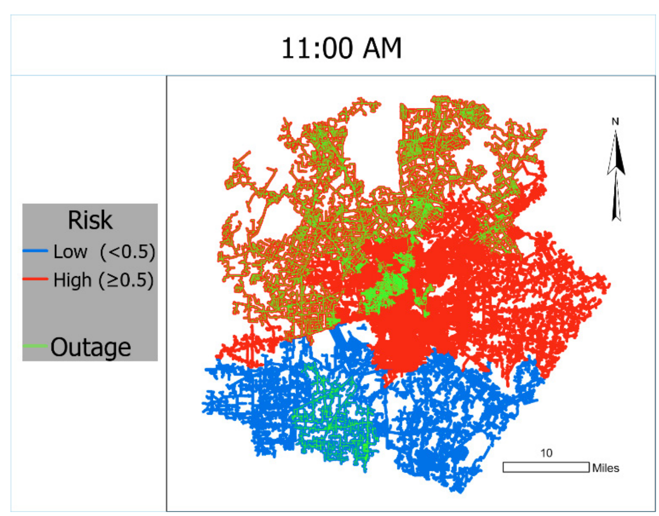

Figure 2. Risk map.

The most used algorithm choices come from extensive Python libraries of ML algorithms (Ramasubramanian and Singh, 2019, Kramer, 2016). The outage SoR prediction implementation indicates that off-the-shelf solutions must be customized to meet the requirements of utility applications. Our experience from outage SoR prediction research supports this hypothesis, as such customized approaches result in improved prediction performance (Baembitov et al., 2024).

2.3 Variability of design options

The variability in the implementation option relates to the several design/implementation choices.

Outage causes, reflected through the related data, determine what types of outages the model predicts. For example, if only weather-related outages are included, the model will focus solely on weather-induced outages.

Spatial and temporal settings define the geographic area and time window for predictions (e.g., citywide vs. districtwide, 1-hour vs. 24-hour predictions).

The prediction window size refers to the period when the predictions are made. In the Figure 3 example, the prediction window is set to five hours, and if an outage occurs in any of those five hours, the prediction window is labeled as 'outage-positive,' or a positive sample. Otherwise, the prediction window is marked as a negative sample. The ML models aim to determine if the window should be classified as positive or negative.

The lead time defines how far the moment of executing the model is from the beginning of the prediction window. This time corresponds to preparing for mitigation actions before an outage. An example of a model run at a two-hour lead time versus zero-hour lead time is illustrated in Figure 3.

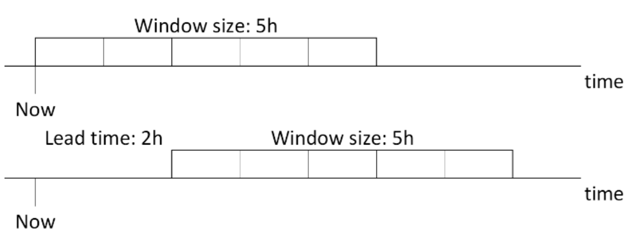

Fig. 3. 0h and 2h lead time for 5h prediction window size.

Prediction frequency is defined as the period between model executions, which may be set at every hour or a longer period.

The input data window refers to the timesteps and all the data connected to those timesteps provided to the model at the execution time. It should not be confused with the prediction window size described above. For example, we can select 6 hours of past weather, 3 hours of forecast data, and 3 weeks of vegetation management records for a single model execution, signifying that the time range for data coming from different sources does not need to be the same.

Dataset selection for training and testing. The examples *are* High-Resolution Rapid Refresh for weather (HRRR) (NOAA, 2024), Light Detection and Ranging (LiDAR) for vegetation and infrastructure mapping (Wanik et al., 2017), and Automated Surface Observing System (ASOS) for meteorological data (NOAA, 2021).

2.4 Interpretability of the results

Developer-focused evaluation metrics. Typically, the metrics used by predictive analytics developers (ex. Accuracy, Precision, Recall, F1 score, Specificity, ROC-AUC, Confusion Matrix) capture different meanings of the results (Powers, 2011). The interpretation is reduced to the choice of metrics used to evaluate the results.

The evaluation metrics results are meaningful to the data analytics experts in assessing the prediction algorithm performance. They are also indispensable for estimating the expected overall efficacy of the model(s), which can be measured by different business metrics (total reduction in equipment downtime, change in customer satisfaction levels, monetary loss reduction, etc.). The business metrics for a given application are derived from performance metrics by considering the cost/reward for each possible outcome: true positive, true negative, false positive, and false negative.

The end-user-focused evaluation. This interpretability goes beyond the data analytics experts as it also concerns the end users (different utility staff and/or electricity consumers). Their interest is in understanding how a prediction for a given outage event should be interpreted to define mitigation measures aimed at reducing the impact of outages. As an example, a possible display for the end user in the utility environment (Fig. 2) illustrates an unfolding set of displays that dynamically shows how the prediction fairs against actual events over time.

The utility personnel have various mitigation options for dealing with outages depending on their responsibility, that is typically constrained by the time interval that they need to react. The control center personnel are issuing network reconfiguration orders to allow alternative power routing to the consumers (if feasible), which happens in the operations time frame of minutes and hours. A set of repair and restoration actions are performed by the outage management crews that can only act in hours, perhaps even days, depending on the environmental conditions. The asset management personnel may use the outage prediction to determine where the outages are more frequent to assess whether some type of equipment needs preventive repair or replacement and if relevant equipment types of replacements are in stock. This last decision-making mitigation measure may take months to implement.

The consumer is in a different position since they can decide the mitigation measures immediately as they receive the outage warning messages, such as rationalizing refrigeration and heating or seeking help in the cooling/warming centers.

This illustrates that the various options in the prediction algorithm design need to be translated to an actionable measure/action that end-users may use to mitigate impacts.

3. INTERPRETABILITY OF THE BASELINE CASE

3.1 The baseline design results

It becomes computationally expensive and, in many cases, infeasible to conduct a sensitivity study of the model performance to all of the design parameters. The need to use many years of data to arrive at reliable estimates exacerbates the problem. Hence, we performed a sensitivity analysis only for the two most important dimensions, spatial resolution, and prediction window size, that correspond to temporal prediction specificity while fixing the rest of the parameters. The selection of the two dimensions was guided by end-user needs to assess the accuracy of the solution at different spatiotemporal resolutions corresponding to a selection of appropriate mitigation measures. We evaluated performance metrics across both dimensions, making it easier to understand how accuracy metrics may change for each baseline case. The baseline results can be viewed using several evaluation metrics, such as F1 score and Average Precision (Table 2), which allows data analytics experts to understand the different performance aspects of the algorithm in question. These results were obtained by the Random Forest algorithm.

Table 2. F1 score || Average Precision results for various number of clusters (1,3,7,86) and different prediction time windows (1,3,5,11)

Window size, h \ N clusters	1	3	7	**86**
1	0.83‖0.78	0.52‖0.48	0.16‖0.22	0.003‖0.02
3	0.95‖0.91	0.83‖0.78	0.46‖0.46	0.018‖0.06
5	0.98‖0.97	0.90‖0.85	0.63‖0.60	0.040‖0.10
11	1.00‖1.00	0.99‖0.97	0.83‖0.79	0.115‖0.19

3.2 Impact of different design implementation approaches

One can strive to improve the baseline at certain intersections of clusters and window sizes by: model choice, data choice, or the choice of hyperparameters directing the model's training. The goal is then to determine how much the baseline case can be improved. This should align with end-user needs, as they would understand specific performance thresholds and their implications for design implementation.

Usually, for ML techniques, extensive models with many parameters are better for modeling a phenomenon (Rizvi, 2023). Improving the outage prediction model to represent the physical nature of how the outage takes place may help the prediction task. Such form of regularization is usually used in Physics Informed Neural Networks (PINN) (Hu et al., 2020).

The values of the baseline case (Table 2) represent the variation over the target cluster's prediction window size and spatial resolution. Increasing the number of clusters and decreasing the window size decreases the prediction accuracy since the number of data cases representing outages is much less for this condition. For larger areas and more considerable window lengths, there are a lot of data points representing outages, which helps train the model better. Frequent prediction, too, can help increase the number of training data points that can be used for training. Larger data can help the models to generalize more effectively for accurate prediction. Variations in prediction lead times can affect accuracy. Predicting outages based on conditions before the event is more accurate when considering the circumstances immediately prior to the event rather than those of the previous day. This approach allows the model to capture better the causality of the external conditions leading to these events, demonstrating that these variables can influence the accuracy of the outage prediction model in different ways.

Various data-based and algorithmic approaches can be explored to improve the prediction models. Recently, prediction of weather parameters has improved, allowing acquiring weather information many days into the future with acceptable accuracy. Moreover, extensive past weather records are available, which can be easily obtained in bulk. Since weather forecasts are easily available and weather explains a lot of occurrences, incorporating the forecast throughout appropriate lead time may result in better outage predictions.

Another way to harness the information provided by weather forecasts is to ingest the forecast parameters (temperature, wind, etc.) in the form of images over the service territory. The multi-input neural network can accommodate several different forecasts at different timestamps inside the prediction window (Chen et al., 2025). This approach allows us to account for several dimensions of the weather forecasts: longitude, latitude, forecast time, and weather parameters.

It is also possible to exclusively focus on the end ranges of spatial and temporal specificity from Table 3, differentiating the predictions across many more clusters and being more specific in the temporal domain. Consequently, more involved and deeper ML models are generally more suitable to tackle such a task w.r.t. model learning capacity. For example, specific model selection can be made by utilizing models resilient to a low number of positive samples in the dataset during training since there are generally significantly more normal operation events than outage events (Rendle, 2010). Aside from using a model robust to the ratio of positive samples, another approach can include adjusting the loss function to accommodate better learning under such a dataset, usually given as the weight of positive or negative samples (Lin et al., 2017).

3.3 The design option impacts on end-use applications

The choice of outage causes driving the input data, and the input window is more apparent to the model developer rather than the consumer. The end-user may be more focused on the spatiotemporal aspects, the prediction window size, along with prediction frequency since such choices help determine when outage risk in a region increases, triggering preventive mitigation actions. Higher predictions frequency keeps utilities informed about sudden changes in SoR, while longer prediction windows improve accuracy, though the exact timing may remain uncertain. A combination of different prediction frequencies and window sizes may allow utilities to better plan mitigation measures. Longer lead times allow more time to prepare, though at the cost of accuracy, whereas shorter windows push utilities to rely on predefined emergency responses.

To understand the impact of the spatial resolution options on customers, one can think of the granularity of the predictions. With a smaller number of clusters formed within the study region, the numerical accuracy of the prediction would be higher. However, it would be difficult for the utility to decide where to implement the countermeasures, harden the infrastructure, or send maintenance crews since the prediction granularity may encompass multiple feeders/substations. It would also require a larger number of consumers to prepare for possible outages. In this regard, the SoR prediction application decision-making end uses might vary depending upon the type of applied mitigation measures and their relative financial costs and benefits.

4. INTERPRETABILITY IMPACT ON DECISION MAKING

4.1 Importance of metrics interpretability

Performance metrics can be misleading when analyzed separately from their intended uses. Each metric aims to condense various factors of outage SoR prediction model performance into a single figure, which inevitably results in some performance information being lost. For this reason, data analytics professionals should use a combination of performance metrics to fully understand the model's predictive capabilities.

Both types of metrics, such as Precision versus Recall and Accuracy and the aggregation of metrics across various temporal and spatial resolutions, should be considered. For example, one utility may be particularly interested in model performance during summer and winter peaks, as disruptions in electricity supplies during these seasons can have significant impacts. A school might concentrate solely on model performance during weekdays and business hours since outages would have a limited effect at other times. An industrial facility within a city limit would be more focused on how the model performs for its specific feeder while being less concerned about city-wide model performance.

Another important consideration is how SoR predictions can improve optimization strategies for outage mitigation. The selection of mitigation measures can pursue different priorities, such as restoring service to the most critical customers (hospitals, emergency facilities) or minimizing overall power quality indices such as SAIFI (System Average Interruption Frequency Index) and SAIDI (System Average Interruption Duration Index) (IEEE, 2012). A utility may prioritize rapid restoration in areas with high economic or societal impact, while another may focus on minimizing the total number of customers affected. Utilities may strike a balance between several objectives by utilizing outage SoRs.

The trustworthiness of the outage prediction task can affect the actions taken by the users to mitigate the effects of outages. Repeated failures to detect outages will discredit the entity generating predictions. Also, several false alarms might impact the level of end-user preparedness. Still, the confidence with which investments are eliminating or mitigating possible outages might be affected by the rate of false alarms, miss rate, and, ultimately, the model's performance. Therefore, it is necessary to enhance model performance with existing and newly available data or enhanced algorithms.

4.2 Importance of end-user decision-making interpretability

After fully understanding the implementation options, the end user of the prediction model must participate in the selection of the design parameters together with the algorithm developers. That way, their needs will be best aligned with resources at their disposal to mitigate the impact.

For instance, power system operators and other facility operators (such as water supply) may prefer outage prediction information with a lead time between a few hours (operations) to a day in advance (operations planning). Upon receiving a real-time SoR feed, an operator may plan and execute grid mitigation and restoration schemes for the affected locations.

If the outage management team has information about the outage 24 hours prior, they might focus on implementing mitigation tasks such as issuing outage alerts, planning a number of repair crews, notifying critical customers, rerouting the supply path, etc. If the prediction is available only a few hours before, the focus will be on preparing the customer response units and restoration steps, such as optimizing the routes for restoration crews to potential outage locations. If the prediction is available several days before, they can make sure that they are well-equipped in terms of manpower as well as spare parts inventory to serve the size of the area to be affected.

In asset management teams' case, if the outage prediction is available months before, they can take proactive steps to harden the infrastructure against the faults and ensure they are prepared once the outage occurs. However, for shorter lead times, they would usually have to wait for the outage to occur and take mitigation steps to remedy the effects of the outage for stakeholders, such as equipment upgrades.

Residential users could benefit from a display that also highlights other essential facilities relevant to their needs in current situations: warming centers during cold weather, water pumping stations supplying their homes, and hospital locations for individuals with chronic health issues.

The selection of design parameters for the algorithm affects the results presentation and ultimately drives the decisions about the mitigation actions available to a given stakeholder. Once these parameters are set and implemented by the algorithm developer, users can expect a certain level of model performance derived from the baseline cases to make informed decisions about deploying mitigation measures in each scenario. Every mitigation action carries a financial cost of implementation but also has potential financial benefits, such as avoided outages, fewer affected customers, and/or reduced restoration time. The final decision on applying mitigation measures depends on balancing these costs and benefits, the level of trust in the model's accuracy (based on historical performance), and the potential consequences of poor outage management (a lawsuit from a major customer, fines, for failing to meet reliability metrics, customer dissatisfaction, and reputational damage). Considering these interpretability factors, a business case can be developed for the implementation of the SoR prediction solution for each individual application use case.

ACKNOWLEDGEMENTS

This work was funded by NSF award # 2125985.

REFERENCES

BAEMBITOV, R., KARMACHARYA, A. M., KEZUNOVIC, M., SARANOVIC, D. & OBRADOVIC, Z. Effect of Lightning Features on Predicting Outages Related to Thunderstorms in Distribution Grids. 58th Hawaii International Conference on System Sciences, Jan. 2025 Waikoloa Village, HI.

BAEMBITOV, R. & KEZUNOVIC, M. 2023. State of Risk Prediction for Management and Mitigation of Vegetation and Weather Caused Outages in Distribution Networks. *IEEE Access,* 11**,** 113864-113875.

BAEMBITOV, R., KEZUNOVIC, M., BREWSTER, K. A. & OBRADOVIC, Z. 2023. Incorporating Wind Modeling Into Electric Grid Outage Risk Prediction and Mitigation Solution. *IEEE Access,* 11**,** 4373-4380.

BAEMBITOV, R., KEZUNOVIC, M. & OBRADOVIC, Z. Graph Embeddings for Outage Prediction. 2021 North American Power Symposium, Nov. 2021. 1-6.

BAEMBITOV, R., KEZUNOVIC, M., SARANOVIC, D. & OBRADOVIC, Z. Sensitivity Analysis of Machine Learning Algorithms for Outage Risk Prediction. 57th Hawaii International Conference on System Sciences, Jan. 2024 Honolulu, HI.

CHEN, S., SAKAI, S., MATSUO-UEDA, M. & UMEMURA, K. 2025. Investigation of Multi-input Convolutional Neural Networks for the Prediction of Particleboard Mechanical Properties. *Scientific Reports,* 15**,** 4162.

DOE 2002-2021. Electric Disturbance Events (OE-417) Annual Summaries. U.S. Department of Energy.

DOKIC, T. & KEZUNOVIC, M. 2019. Predictive Risk Management for Dynamic Tree Trimming Scheduling for Distribution Networks. *IEEE Transactions on Smart Grid,* 10**,** 4776-4785.

DUNN, S., WILKINSON, S., ALDERSON, D., FOWLER, H. & GALASSO, C. 2018. Fragility Curves for Assessing the Resilience of Electricity Networks Constructed from an Extensive Fault Database. *Natural Hazards Review,* 19**,** 04017019.

GARLAND, J., BAKER, K. & LIVNEH, B. Weather-induced Power Outage Prediction: A Comparison of Machine Learning Models. 2023 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids, Oct.-Nov. 2023. 1-6.

HU, X., HU, H., VERMA, S. & ZHANG, Z.-L. 2020. Physics-guided Deep Neural Networks for Power Flow Analysis. *IEEE Transactions on Power Systems,* 36**,** 2082-2092.

IEEE 2012. IEEE Std 1366 Guide for Electric Power Distribution Reliability Indices.

JAIN, A., SHAH, T., YOUSEFHUSSIEN, M. & PANDEY, A. Combining Remotely Sensed Imagery with Survival Models for Outage Risk Estimation of the Power Grid. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, June 2021. 1202-1211.

KEZUNOVIC, M., XIE, L. & GRIJALVA, S. The Role of Big Data in Improving Power System Operation and Protection. 2013 IREP Symposium Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid, Aug. 2013. 1-9.

KHOSHJAHAN, M., BAEMBITOV, R. & KEZUNOVIC, M. 2023. Informed Prosumer Aggregator Bidding Strategy via Incorporating Distribution Grid Outage Risk Predictions. *IEEE Access,* 11**,** 28585-28595.

KOR, Y., REFORMAT, M. Z. & MUSILEK, P. Predicting Weather-related Power Outages in Distribution Grid. 2020 IEEE Power & Energy Society General Meeting, Aug. 2020. 1-5.

KRAMER, O. 2016. Scikit-Learn. *Machine Learning for Evolution Strategies.* Cham: Springer International Publishing.

LIN, T. Y., GOYAL, P., GIRSHICK, R., HE, K. & DOLLÁR, P. Focal Loss for Dense Object Detection. 2017 IEEE International Conference on Computer Vision, Oct. 2017 2017. 2999-3007.

MUJJUNI, F., BETTS, T. R. & BLANCHARD, R. E. 2023. Evaluation of Power Systems Resilience to Extreme Weather Events: A Review of Methods and Assumptions. *IEEE Access,* 11**,** 87279-87296.

NOAA. 2021. *Automated Surface Observing Systems* [Online]. NOAA's National Weather Service. Available: https://www.weather.gov/asos/asostech [Accessed Jan. 2021].

NOAA. 2024. *High-Resolution Rapid Refresh (HRRR) Model* [Online]. Available: https://registry.opendata.aws/noaa-hrrr-pds [Accessed Jan. 2024].

OWERKO, D., GAMA, F. & RIBEIRO, A. Predicting Power Outages Using Graph Neural Networks. 2018 IEEE Global Conference on Signal and Information Processing, Nov. 2018. 743-747.

POWERS, D. M. W. 2011. Evaluation: From Precision, Recall and R-measure to ROC., Informedness, Markedness & Correlation. *Journal of Machine Learning Technologies,* 2**,** 37-63.

RAMASUBRAMANIAN, K. & SINGH, A. 2019. Deep Learning Using Keras and TensorFlow. *Machine Learning Using R: With Time Series and Industry-Based Use Cases in R.* Berkeley, CA: Apress.

RENDLE, S. Factorization Machines. 2010 IEEE International Conference on Data Mining, Dec. 2010 2010. 995-1000.

RIZVI, M. 2023. Leveraging Deep Learning Algorithms for Predicting Power Outages and Detecting Faults: A Review. *Advances in Research,* 24**,** 80-88.

UDEH, K., WANIK, D. W., CERRAI, D., AGUIAR, D. & ANAGNOSTOU, E. Autoregressive Modeling of Utility Customer Outages with Deep Neural Networks. 2022 IEEE 12th Annual Computing and Communication Workshop and Conference, Jan 2022. IEEE, 0406-0414.

WANG, X., FATEHI, N., WANG, C. & NAZARI, M. H. 2024. Deep Learning-based Weather-related Power Outage Prediction with Socio-economic and Power Infrastructure Data. *arXiv preprint arXiv:2404.03115*.

WANIK, D. W., PARENT, J. R., ANAGNOSTOU, E. N. & HARTMAN, B. M. 2017. Using Vegetation Management and LiDAR-derived Tree Height Data to Improve Outage Predictions for Electric Utilities. *Electric Power Systems Research,* 146**,** 236-245.