

Received October 29, 2021, accepted November 8, 2021, date of publication November 9, 2021, date of current version November 12, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3127042

Detecting and Interpreting Faults in Vulnerable Power Grids With Machine Learning

ODIN FOLDVIK EIKELAND¹, INGA SETSÅ HOLMSTRAND², SIGURD BAKKEJORD²,
MATTEO CHIESA¹, AND FILIPPO MARIA BIANCHI^{1,3,4}

¹Department of Physics and Technology, UiT-The Arctic University of Norway, 9019 Tromsø, Norway

²Arva Power Company

³Department of Mathematics and Statistics, UiT-The Arctic University of Norway, 9019 Tromsø, Norway

⁴NORCE The Norwegian Research Centre AS, 9294 Tromsø, Norway

Corresponding author: Filippo Maria Bianchi (filippo.m.bianchi@uit.no)

The work of Odin Foldvik Eikeland, Matteo Chiesa, and Filippo Maria Bianchi was supported in part by the Project “Transformation to a Renewable & Smart Rural Power System Community (RENEW)” under Grant 310026, and in part by the Arctic Centre for Sustainable Energy (ARC), UiT-The Arctic University of Norway.

ABSTRACT Unscheduled power disturbances cause severe consequences both for customers and grid operators. To defend against such events, it is necessary to identify the causes of interruptions in the power distribution network. In this work, we focus on the power grid of a Norwegian community in the Arctic that experiences several faults whose sources are unknown. First, we construct a data set consisting of relevant meteorological data and information about the current power quality logged by power-quality meters. Then, we adopt machine-learning techniques to predict the occurrence of faults. Experimental results show that both linear and non-linear classifiers achieve good classification performance. This indicates that the considered power quality and weather variables explain well the power disturbances. Interpreting the decision process of the classifiers provides valuable insights to understand the main causes of disturbances. Traditional features selection methods can only indicate which are the variables that, on average, mostly explain the fault occurrences in the dataset. Besides providing such a global interpretation, it is also important to identify the specific set of variables that explain each individual fault. To address this challenge, we adopt a recent technique to interpret the decision process of a deep learning model, called Integrated Gradients. The proposed approach allows gaining detailed insights on the occurrence of a specific fault, which are valuable for the distribution system operators to implement strategies to prevent and mitigate power disturbances.

INDEX TERMS Energy analytics, machine learning interpretability, power quality disturbances.

I. INTRODUCTION

Unscheduled power disturbances cause problems for customers and grid operators as they affect all customers connected to the power network, from single households to large industries [1]–[4]. Power failures might have complex and adverse socio-economic consequences in communities that are heavily reliant on the electricity supply [5], [6]. The distribution system operator (DSO) is contractually obliged to provide a reliable power supply and to compensate customers affected by power interruptions [7]. To meet the expected energy demand, the DSOs must implement management plans that account for the underlying infrastructure.

In this study, we focus on disturbances on a power grid in an Arctic region in Northern Norway, where the energy demand from local food industries has increased greatly. The growth in energy demand has resulted in more frequent power disturbances, as the current power grid is operating close to its maximum capacity. One way to improve the reliability of the power supply is to build a new distribution grid that can handle larger power demand. However, this is costly, time-consuming, has a huge environmental impact, and contradicts the vision of better utilizing the current electricity grid infrastructure* [8]. An alternative solution is to limit the failures and strengthen only the most vulnerable parts of the grid, but this requires first identifying the factors that trigger power disturbances.

The associate editor coordinating the review of this manuscript and approving it for publication was Nagesh Prabhu¹.

*<https://www.miljodirektoratet.no/publikasjoner/2020/januar-2020/klimakur2030/>

The identification of causing factors of faults in the power grid has proven to be a major challenge for the DSO [4]. However, the increased availability of energy-related data makes it possible to exploit advanced data analytics techniques to support the development of strategies for improving the reliability of the power grid [9]–[15]. Recent studies based on statistical data analysis and machine learning (ML), indicated that extreme weather conditions are often an important cause of faults in power grids [16]–[21]. However, other factors besides weather could likely affect the power quality.

In this work, we explore a wide spectrum of potential causing factors for power failures. We consider explanatory variables relative both to weather and high-resolution power-quality data. We adopt ML techniques to detect the power disturbances and to identify the factors that mostly explain the power disturbances.

This paper extends our previous study, which analyzed fault data in the Arctic power grid during the year 2020 [22]. There were important shortcomings in the data used in our previous work:

- 1) The machines of the local industries connected to the power grid are so sensitive to the power quality that they experience failures that are not registered in the failure-reporting system of the DSO.
- 2) The resolution of data in 2020 was too low (1-hour) to understand how power consumption truly affects power quality.

To address these issues, new power quality meters were installed on 19 February 2021 in the power grid under analysis. These meters log data every minute and register every small voltage variation. In addition, they provide detailed information about the power quality in the grid, such as the specific phase where the fault is registered, the magnitude of voltage variation, frequency imbalance, and the amount of flicker.

Contributions: First, we build a power faults classification dataset in collaboration with the DSO, by collecting variables that are considered as most relevant in explaining power disturbances. Then, we train different classifiers, including linear classifiers and a deep learning architecture, to detect an incoming fault from the weather and power-quality variables, registered one minute before the specific fault occurs. As shown in the experimental results, the classifiers manage to detect most of the power disturbances before their onset, demonstrating that high-resolution data from power quality meters in conjunction with weather data are highly informative.

To gain a better understanding of the relationships between the different variables and the power disturbances, we analyze the decision process of the classifiers. First, we consider traditional features selection methods, which identify which are the most important variables in the dataset that explain the fault occurrence. While such an approach gives a global overview of the variables that are, on average, the most informative in the dataset, it does not allow to reason about specific cases.

To address this challenge, we adopt a recent technique to interpret the decision process of a deep learning model, called Integrated Gradients (IG). For each individual sample, IG assigns to each feature a score, whose magnitude indicates how much the value of such feature contributes to determine the class of the sample. The proposed methodology shows that the classifiers focus on heterogeneous sets of features when processing different samples. This indicates that the occurrence of faults can be explained by multiple different patterns in the weather and power-quality variables. Our findings are valuable to the DSO for implementing strategies to prevent and mitigate power disturbances.

II. RELATED WORK AND STUDIES

There exist a vast amount of literature about the detection of different classes of power quality disturbances, such as deviation in voltage, current, and frequency signals. For example, Ref. [23] provides a comprehensive review of more than 150 research studies between 1986 and 2014 on detection and classification of power quality disturbances. In another comprehensive and more recent survey, [24] reviewed 242 papers on Power Quality and Classification (PQD&C) techniques based on digital signal processing and ML. The survey performed a comparative assessment on various PQD&C techniques by considering several criteria, such as type of data used, type of PQ disturbance, and classification accuracy.

However, fault detection and classification is a reactive process where models try to classify the fault after it has occurred. On the other hand, it is often interesting to identify the causing factors and predict the onset of a power fault. A fault prediction model should be able to quantify the likelihood of observing a fault in the next period given a set of conditions described by the explanatory variables in the model. Additionally, the identification of causing factors for faults will help the DSO to implement strategies to prevent and mitigate incoming faults.

There exist some prior relevant work on identifying causing factors for faults in the power grid. The causing factors are often divided into two different categories: i) weather conditions, and ii) other factors such as human-related activities (energy consumption).

A. WEATHER-RELATED FAULTS

Harsh and severe weather events are considered to be an important source of faults, and several studies have been conducted to address the impact of such events on power quality.

Owerko *et al.* predicted power faults in New York City by monitoring weather conditions [21]. The authors deployed a Graph Neural Network to model the spatial relationships between weather stations and improve the prediction performance.

The impact of seasonal weather on forecasting power disturbances was investigated in [25]. The authors tested the performance of the proposed models by using two different training sets: seasonal or all-year data. It was shown that,

in some cases, the prediction performance of the models improved when the training data is limited to a subset corresponding to a particular meteorological season.

The impact of weather variations and extreme weather events on the resilience of energy systems was investigated in [16]. The authors developed a stochastic-robust optimization method to consider both low impact variations and extreme events. The method was applied on 30 cities in Sweden. The results indicated that 16% drop in power supply reliability is due to extreme weather events.

Other examples of relevant work on weather-related faults can be found in Refs. [17]–[20]. In addition, several risk assessment studies on the impacts of extreme weather hazards such as earthquakes, thunderstorms, and hurricanes can be found in Refs. [26]–[32].

The works mentioned so far consider only severe weather events and disregard other factors, such as heavy energy load caused by human-related activities. Additionally, many methodologies are tested on synthetic data or on public benchmark datasets, which limits the scope of the evaluation and poses constraints on the data acquisition procedure.

B. ALTERNATIVE APPROACHES FOR FAULT DETECTION

A methodology to predict power faults by analyzing advanced measurement equipment such as Power Quality Analyzers (PQAs) and Phasor Measurement Units (PMUs) has been proposed in [33]. The study used real-world measurements from nine PQA nodes in the Norwegian grid to predict incipient interruptions, voltage dips, and earth faults. The authors find incipient interruptions easiest to predict, while earth faults and voltage dips are more challenging to predict.

The authors in [34], compared several ML methods to predict power disturbance events such as voltage dips, ground faults, rapid voltage changes, and power interruptions. The Random Forest models achieved the highest performance and the results indicated that voltage dips and rapid voltage changes were the easiest to predict.

The challenge of detecting back-fed ground-faults has been recently addressed in [35]. The authors show that faults can be detected by integrating advanced metering infrastructure with a distribution management system. However, the proposed solution is relevant only for DSOs that adopt the OpenDSS software.

The study in [36] investigated the possibility of predicting voltage anomalies minutes in advance by using an ML model trained on historical power quality analyzers (PQA) data. The voltage data were collected from 49 measuring locations in the Norwegian power grid. The model attempted to predict voltage anomalies 10 minutes in advance based on the presence of early warning signs in the preceding 50 minutes. It was found that the time passed since the previous fault is a major factor that affects the probability of a new imminent fault.

In [37], the application of clustering and dimensionality reductions techniques to predict unscheduled events

was investigated. First, the authors used several techniques to reduce the dimensionality of the data and to cluster events based on analytical features. Then, the fault events were separated from the normal operating conditions. The findings show promising results when using balanced datasets, while the predictive capability is significantly reduced in unbalanced datasets that, however, often appear in real-world case studies.

Other relevant work on fault detection based on ML techniques can be found in Refs. [38]–[43]. In addition, there is some relevant work that adopts novel ML techniques for detecting and localizing faults in the power distribution network [9]–[12].

This section presented several relevant works in predicting faults by assessing either weather effects or human activities. One of the goals of our work is to consider, at the same time, a larger amount of weather variables and electricity-related measures as potential causes of power disturbances. A close collaboration with the local DSO has provided us with valuable insights about the relevant variables that should be monitored to construct a new classification dataset. More importantly, none of the previous work has focused on interpreting the decision process of the classifier, which is key to understanding the causes of faults and can provide valuable information to improve the power grid reliability.

III. POWER FAULTS DATASET

In this study, we focus on a power grid with a radial structure located in the Arctic. A detailed description of the grid configuration is deferred to Sect. A in the Appendix. The grid is subject to frequent power faults, which could be caused by weather factors or by the strain of the infrastructure from a local industry, which dominates the load consumption in the power grid.

We prepared a classification dataset where each sample refers to a period when the grid is operating in normal conditions or to a period preceding a fault, respectively. Each sample is associated with a feature vector $x \in \mathbb{R}^{12}$ and a label $y \in \{0, 1\}$, indicating the normal condition or the imminent fault, respectively. The feature vector consists of 6 different energy-related variables and 6 different weather variables, summarized in Tab. 1. A fault is registered when there is at least a 10% drop in voltage magnitude. Further details about faults measurement, what the weather and power variables represent, and how they are collected, are described in Sect. B in the Appendix.

The dataset contains 90 samples representing reported faults ($y = 1$), which occurred in the period between 19.02.2021 to 30.04.2021. Naturally, the amount of samples associated with normal operating conditions is much larger. In addition, in normal operating conditions the values x from neighboring hours are very similar to each other. To limit the amount of class imbalance in the dataset and the redundancy in the over-represented class, we arbitrarily subsampled the non-fault class ($y = 0$) by taking 1 sample every 60. In the final dataset, there are 90 samples representing a reported

TABLE 1. Variables analyzed to detect faults in the power grid.

Feature ID	Weather variables
1	Wind speed of gust
2	Wind direction
3	Temperature
4	Pressure
5	Humidity
6	Precipitation
Power variables	
7	Difference in Frequency
8	Difference in Voltage imbalance
9	Difference in Active Power
10	Minimum Power Factor
11	Difference in Reactive Power
12	Flicker

fault and 1, 647 samples representing normal operating conditions without any power disturbance.

IV. METHODOLOGY

Our approach consists of two steps. First, we train a classifier to predict the onset of power faults given the value of the electricity and weather variables. If we obtain a high classification accuracy, we can conclude that there are strong relationships between the variables, x , and the occurrence of faults, y . Then, we use two different techniques to highlight the most informative features identified by the classifiers to solve the task.

In Sect. IV-A and IV-B, we describe which classifiers are considered in this study. In Sect. IV-C, we present an approach for interpreting the decision process of a neural network classifier.

A. LINEAR CLASSIFIERS

We consider three different linear classifiers. The first is a Ridge regression classifier, which first converts the target values into $\{1, 1\}$ and then treats the problem as a regression task [44]. The second model is Logistic regression, which uses a logistic function to approximate the probability of binary classification variable [44]. The third model is the Linear Support Vector Classification model (LinearSVC), which is a type of Support Vector Machine (SVM) [45] endowed with a linear kernel.

Due to the strong class imbalance, we configure each model to assign a class weight that is inversely proportional to the number of samples in each class. In this way, errors on the underrepresented class (faults, $y = 1$) are penalized much more than errors on the larger class (nominal condition, $y = 0$).

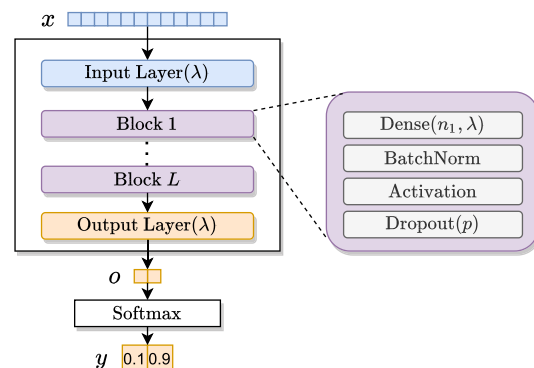
One advantage of using linear classifiers is that they construct a decision boundary directly in the input space, which allows to easily interpreting the decision process of the classifier. In particular, the linear models assign a weight w_i to each feature x_i in the input space: the higher w_i , the more the values of x_i impact the classification outcome. Therefore, looking at the magnitude of the weights w_i is a simple strategy to

estimate the average importance of the features in the dataset for the classification task.

B. NON-LINEAR CLASSIFIERS

We consider two non-linear classifiers. The first is a non-linear SVC classifier equipped with a radial basis function kernel (RBF SVC). As for the linear models, also in this case we used class weights inversely proportional to the class frequency.

The second non-linear classifier considered is a Multi-Layer Perceptron (MLP) [46]. The MLP consists of an input layer that takes input vectors $x \in \mathbb{R}^{12}$, L hidden blocks, an output layer that generates a 2-dimensional output $o \in \mathbb{R}^2$, and a softmax activation that gives the vector of class probabilities y . Each block l consists of a dense layer with n_l units, a Batch Normalization layer [47], a non-linear activation function, and a Dropout layer [48] with dropout probability p . All trainable weights in the MLP, except the biases, are regularized with L_2 -norm penalty with strength λ . Fig. 1 depicts the MLP architecture.

**FIGURE 1. Architecture of the MLP.**

The MLP is trained by minimizing a cross-entropy loss, using batches of size b , and the Adam optimization algorithm [49] with initial learning rate r . Due to the strong class imbalance in the dataset, we initially trained the MLP by weighting the loss of each sample with a value inversely proportional to the class frequency, as we did for the other classifiers. However, we found out that the MLP achieved better performance by re-sampling the minority class during training. This allows achieving class balance at the expense of introducing redundancy, by re-proposing the same samples multiple times. We also tried to achieve class balance by subsampling the majority class but, due to the small number of samples in the fault class, the total number of inputs in each training epoch was too small and the samples from the majority class were shown too few times during training.

C. INTERPRETATION OF THE MLP RESULTS WITH INTEGRATED GRADIENTS

In the following, we introduce the technique adopted to interpret the decision process of the MLP. A short review of important approaches for interpretability in deep learning,

which have been proposed over the past few years (and briefly mentioned hereafter), is deferred to Sect. C in the Appendix.

Integrated Gradients (IG) [50] is a technique proposed to satisfy two axioms, which are not jointly enforced by other existing attribution schemes (see Sect. C for details). According to the first axiom, *sensitivity*, if the input and an *uninformative baseline* differ in exactly one feature, such a feature should be given non-zero attribution. While interpretability approaches such as LRP [51] and DeepLiFT [52] ensure sensitivity due to the conservation of total relevance, gradient-based methods [53]–[56] do not guarantee the sensitivity axiom because the saturation at ReLU or MaxPool makes the score function locally “flat” with respect to some input features.

The second axiom, *implementation invariance*, states that when two models are functionally equivalent, they must have identical attributions to input features. While implementation invariance is mathematically guaranteed by vanilla gradient approaches, the coarse approximation to gradients in LRP and DeepLiFT might break this assumption.

The attribution to feature i given by IG is

$$\text{IG}_i(x) ::= (x_i - x'_i) \times \int_{\alpha=0}^1 \frac{\partial F(x' + \alpha \times (x - x'))}{\partial x_i} d\alpha, \quad (1)$$

where i is an input feature, x is a sample in the dataset, x' is the uninformative baseline, and α is an interpolation constant used to perturb the features of the input sample. The above definition ensures both the desirable assumptions:

- By the Fundamental Theorem of Calculus, IGs sum up to the difference in feature scores and, thus, follow sensitivity;
- Since the IG attribution is completely defined in terms of gradients, it ensures implementation invariance.

IG has become a popular interpretability technique due to its broad applicability to any differentiable neural network model, ease of implementation, theoretical justifications, and computational efficiency.

Implementation: IG is a post-hoc explanatory technique that works with any differentiable model, $F(\cdot)$, regardless of its implementation. In this paper, we let $F(\cdot)$ be the MLP model described in Section IV-B that takes as input tensor the feature vector $x \in \mathbb{R}^{12}$ and generates an output prediction tensor, $o = F(x)$, called *logit*. In our case, $o \in \mathbb{R}^2$ and $\text{softmax}(o)$ gives the probability of x being “fault” and “non-fault”.

The baseline x' in (1) is an uninformative input used as a starting point to compute the IG attributions. The baseline is essential to interpret the IG attributions as a function of individual input features. It is important to choose a baseline that encodes as much as possible the lack of information about the target class c . In a classification task with multiple classes, we want $\text{softmax}[F(x')]_c \approx 0$. In a binary classification task, like in our case, we can choose a baseline that gives equal probability of belonging to both classes, i.e., $\text{softmax}[F(x')]_0 \approx \text{softmax}[F(x')]_1 \approx 0.5$. In computer vision tasks, a black image (all pixels at 0) is




FIGURE 2. Class probabilities for different baselines on the power-faults dataset.

commonly used as a baseline. However, in our dataset, the value 0 might actually be informative because the absence of some specific features can increase the probability of belonging to a specific class (e.g., in the absence of wind it is less likely to observe a fault). Fig. 2(left) shows that the MLP assigns with high confidence the zero-baseline x'_z to class 0 (non-fault). Therefore, different alternatives should be considered as the baseline. One option is to cast the binary classification problem into a 3-classes problem and re-train the two that assign a vector of zeros to a third, dummy class. In this way, when using the zero-baseline x'_z , we would get $\text{softmax}[F(x'_z)]_0 \approx \text{softmax}[F(x'_z)]_1 \approx 0$. Other alternatives are to use a mean-baseline, x'_m , which is a vector computed as a weighted average of the features across the two classes or to use, or a random baseline x'_r . In the latter case, the final result is given by averaging the IG attributions computed from several random baselines. As shown in Figure 2, the mean baseline gives almost the same probability to classes 0 and 1, while the random baseline has the tendency to assign a strong probability to one of the two classes. Therefore, we used the mean baseline in all our experiments.

The default path used by the integral in (1) is a straight line in the feature space from baseline to the actual input. Since the choice of path is inconsequential with respect to the above axioms, we use the straight-line path that has the desirable property of being symmetric with respect to both x and x' . The numerical computation of a definite integral is often not tractable and is necessary to resort to numerical approximations. The Riemann trapezoidal sums offer a good trade-off between accuracy and convergence and changes (1) into:

$$\text{IG}_i^{\text{approx}}(x) ::= (x_i - x'_i) \times \sum_{k=1}^m \frac{\partial F(x' + \frac{k}{m} \times (x - x'))}{\partial x_i} \times \frac{1}{m}, \quad (2)$$

where m is the number of finite steps used to approximate the integral and $\alpha \approx k/m$. The m samples $\mathcal{X} = \{x' + \frac{k}{m} \times (x - x')\}_{k=1}^m$ represent the linear interpolation between the baseline and the input. Fig. 3 depicts such an interpolation path from the mean-baseline to a specific sample of class “fault” in our dataset.

After generating the set of interpolated samples \mathcal{X} , we can compute the gradients $\frac{\partial F(\mathcal{X})}{\partial x_i}$ that quantify the relationship

between the changes in the input features and the changes in the predictions of the MLP F . Important features will have gradients with steep local slopes with respect to the probability predicted by the model for the target class. Interestingly, the largest gradient magnitudes generally occur during the first interpolation steps. This happens because the neural network can saturate, meaning that the magnitude of the local feature gradients can become extremely small and go toward zero resulting in important features having a small gradient. Saturation can result in discontinuous feature importances and important features can be missed. This is the key motivation why rather than simply using the gradients of the actual input, $\frac{\partial F(\mathcal{X})}{\partial x_i}$, IG sums all the gradients accumulated during the whole interpolation path. This concept is exemplified in Fig. 4(left), showing that the model prediction quickly converges to the correct class in the beginning and then flattens out. There could still be less relevant features that the model relies on for correct prediction that differs from the baseline, but the magnitudes of those feature gradients become really small, as shown in Fig. 4(right). The Figure is obtained using the same data of Fig. 3.

V. EXPERIMENTAL EVALUATION

After introducing the experimental setting, in Sect. V-A we compare the classification performance of the different classifiers on our dataset. Then, in Sect. V-B we first analyze the specific samples of class “fault” that are missed by the classifiers and, then, we consider two techniques to interpret the decision process of the classifiers.

Model Selection and Performance Evaluation: The linear and the SVM classifiers are implemented with the scikit-learn library,[†] while the MLP is implemented in Tensorflow.[‡] To evaluate the model performance we first shuffle the data and then perform a stratified k-fold, with $k = 5$. In each fold, 80% of the data are used as a training set, and the remaining 20% is used as a test set. The training is further divided into two parts: 80% is used to fit the model’s coefficients and 20% is used as a validation set to find the optimal hyperparameters.

The hyperparameters of the linear models and the SVM are optimized with a grid search. In particular, we optimize the regularization strength in the Ridge regression classifier, Logistic regression, and LinearSVC. For the non-linear SVM classifier, we also optimize the width of the radial basis function.

For the MLP, due to the higher amount of hyperparameters and the longer training time, we used the Bayesian optimization strategy implemented in Keras Tuner[§] and evaluated a total of 5,000 configurations. In particular, we optimized the number of layers L , the number of units n_l in each layer, the L_2 regularization coefficient λ , the dropout probability p , the learning rate r , and the type of activation function (ReLU, tanh, or ELU). We used a fixed batch size $b = 32$, an early

stopping with patience of 30 epochs, and we reduced the initial learning rate by a factor of 1/2 when the validation loss was not decreasing for 10 epochs.

Before training the models, the input values x are normalized feature-wise by subtracting the mean and dividing by the standard deviation computed on the training set. The overall performance of each classification model is the average performance obtained on each test set of the 5 folds.

Performance Measures: The classification performance is measured by looking at the confusion matrix, which reports the following quantities: True Negatives (TN) – correctly identified non-faults, False Positives (FP) – non-faults predicted as faults, False Negatives (FN) – faults missed, and True Positives (TP) – faults correctly identified. To quantify the performance with a single value we use the F1 score:

$$F1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} = \frac{TP}{TP + \frac{FP+FN}{2}}. \quad (3)$$

Due to the strong class imbalance in the dataset, we compute a weighted F1 score, i.e., we weight the F1-score obtained for each class by the number of samples in that class and then we compute the average:

$$F1_{\text{weighted}} = \frac{(n_{\text{faults}} \cdot F1_{\text{faults}}) + (n_{\text{non-faults}} \cdot F1_{\text{non-faults}})}{n_{\text{faults}} + n_{\text{non-faults}}}, \quad (4)$$

where $n_{_}$ and $F1_{_}$ indicate the number of samples and classification scores for each class, respectively.

Selecting the Number of Interpolation Steps in IG: The result of the IG attribution depends on the number of steps m (see Eq. 2). One of the properties of IG is completeness, meaning that feature attributions encompass the entire prediction of the model. As a consequence, the importance score should capture the individual contribution of each feature to the prediction. Therefore, by adding together all the importance scores is it possible to recover the entire prediction value for a given sample x . In particular, we have that the variation in classification score (e.g., the probability of being a fault) is

$$\delta = \sum_i IG_i(x) - (F(x)_c - F(x')_c)$$

where $F(x)_c$ and $F(x')_c$ are the prediction scores for class c when the model takes as input x and x' , respectively. Since we want the $\sum_i IG_i(x)$ to explain the whole difference in the class attributions, the number of integration steps m should be increased until when δ becomes as close as possible to zero. Following this principle, we found $m = 100$ to be sufficiently large for our experiments as it gives $\delta < 1e - 2$.

A. CLASSIFICATION PERFORMANCE OF THE DIFFERENT METHODS

Here, we compare the classification performance obtained by the linear methods, SVM classifier, and the MLP. The classification performance of each model is reported in Tab. 2 in terms of average Weighted F1 score and the average number

[†]<https://scikit-learn.org/>

[‡]<https://www.tensorflow.org/>

[§]https://keras.io/keras_tuner/




FIGURE 3. Top row: linear interpolation from the mean-baseline (left) to an actual sample of class fault (right). Bottom row: classification probabilities assigned by the MLP at each step of the interpolation.




FIGURE 4. Left: target class predicted probability over α . Right: average feature gradients (normalized) over α .

TABLE 2. Classification score for different models.

	Classifier	TN	FP	FN	TP	Weighted F1 score
Linear	Ridge Classifier	272	57	4	14	0.785
	Logistic regression	275	54	5	13	0.756
	LinearSVC	276	54	5	13	0.757
Non-linear	RBFSVC	283	46	5	13	0.771
	MLP	285	45	4	14	0.803

of TN, FP, FN, and TP obtained across the 5 folds. Note that the TN, FP, FN, and TP are rounded to the closest integer.

The MLP classifier achieves top performance with a weighted F1 score of 0.803, followed by the Ridge Classifier and the SVC with RBF kernel that achieves weighted F1 scores 0.785 and 0.771, respectively. In our case study, is important to miss as few faults as possible, meaning that solutions with very few FN (missed detection) are acceptable even if the number of FP (false alarms) is significant. The MLP and Ridge Classifier provides the most promising result with 4 FN and 14 TP on average.

Finally, it is interesting to notice that linear and non-linear models achieve similar performance. This suggests that the two classes are almost linearly separable, i.e., most of the

data samples can be separated reasonably well by an hyper-plane in the input features space. On the other hand, the data samples that are misclassified are very entangled, and is difficult to find a decision boundary, even if is non-linear, that can correctly separate them. The good performance of the classification models motivates the feature interpretation procedure discussed in the next section.

B. ANALYSIS AND INTERPRETATION OF THE RESULTS

For the next analysis, we generate a fixed random train/validation/test split and used the same fold for each model. This allows us to analyze in detail the solution obtained by the different methods on a single test set, which contains 18 faults and 330 non-faults. Interestingly, all models fail to correctly classify as faults the same 5 data samples. A closer manual investigation on such 5 samples shows the following:

- 1) **2021-02-22 at 19:29:00**: is an empty measurement,
- 2) **2021-02-22 at 21:55:00**: is a phase-to-ground fault;
- 3) **2021-02-22 at 22:12:00**: is a phase-to-ground fault;
- 4) **2021-02-26 at 11:58:00**: is an actual fault that was missed by the classifiers;
- 5) **2021-03-02 at 09:29:00**: is a fault with an unusual long duration.

The first FN could have been caused by some type of error, such as a calibration error, in the measurement instruments.

In the case of a ground fault, the electrical transformers connected to the grid break, and the power that flows through the transformer flows to the ground. When the end of the electrical transformer station that contacts the ground level is on the downstream side, a ground fault occurs [35]. The ground fault is detected as a reduction of only one of the three-phase voltages. Fig. 5 depicts the phase voltages when the first ground fault occurred: it is possible to see that




FIGURE 5. The phases in a phase-to-ground fault incident. The ground fault occurs on Phase A which is decreasing significantly (voltage drop), while the other two in the three-phase system have a minor decrease from nominal voltage level.

Phase A decreases significantly, while the other two stay above the nominal voltage value. It is difficult to explicitly detect ground faults from only weather and electricity load measures considered as input variables, and therefore it is reasonable that the models miss the faults number 2 and 3.

Similar to the ground faults, the 4th FN could be caused by a factor not described in the weather and electricity variables. For example, it could have been caused by vegetation or animals interacting with the power lines.

Finally, the 5th FN is a fault that lasts for 200 seconds, while the usual duration of the faults is approximately 25-30 secs. This suggests that the fault is an anomaly that is not well represented in the dataset and, therefore, is difficult to be classified accurately.

To identify the most important variables that explain the faults, we try to interpret the decision process of the different models. First, we analyze the coefficients of the linear models, which give a “global” interpretation of the variables importance. Then, we use the IG technique for a “local” interpretability of the features that explain the class of a specific data sample.

1) GLOBAL INTERPRETABILITY

As discussed in Sect. IV-A, when using linear models we can interpret the magnitude of the weights assigned to the input features as the global importance of the features for the classification problem. Fig. 6 reports the feature weights learned by the three different classifiers. We observe that in each model the *Wind speed of gust* variable is always associated with a weight with a large magnitude. The Linear SVC and the Logistic Regression classifiers attribute large importance also to the *Flicker* variable, while the Ridge Regression classifiers weight the other features more uniformly and assign weights to *Temperature* and *Humidity* that are slightly larger than the weight assigned to *Flicker*.

This analysis suggests that both the industrial activity and the weather effects are important in discriminating between the fault and non-fault classes. According to the linear models, the most important among the power-related variables seems to be *Flicker*, while the *Wind speed of gust* is consistently the most explanatory weather-related variable.




FIGURE 6. Coefficients’ magnitude assigned to each feature by different linear models. High magnitude indicates that the corresponding feature is important.

These observations are aligned with the experiences of the DSO and the local customers, as more faults seem to occur when there is high activity at the industries and the machines operates at full load. In addition, it has been noted that faults are more likely to occur when there is a strong wind, which could create collisions in the cables of the power lines.

2) LOCAL INTERPRETABILITY

The faults correctly classified by the different models are reported in Tab. 3, together with the confidence score of the MLP classifier. The confidence score can be interpreted as the probability that the MLP believes a sample is a fault. The MLP correctly classifies with high confidence most fault samples and assigns a probability greater than 90% to 5 out of 13 samples. As a side note, the faults do not appear to be clustered around specific days or periods, but they seem to be uniformly distributed over time.

We focus on the samples 52, 140, 227, 304, and 316 in Tab. 3, which are those classified with the highest confidence, and we use IG to identify which are the variables that are most important for the MLP to determine the correct fault class. The results are reported in Fig. 7. The top-left plot depicts the uninformative baseline, which corresponds to what an “average” sample in the dataset looks like. The blue bar plots represent the value of the 12 features in the 5 selected samples. Finally, the green and red bar plots are the output of the IG procedure.

The green bars indicate that a feature is important for the classification result. The higher the i -th green bar, the more the feature value x_i in the sample (blue bar) explains the classification result, compared to the value x_i' in the baseline (black bar). For example, in Sample 227, the value of *Flicker* is much greater than in the baseline. IG assigns a high score (tall green bar) to this difference, meaning that the MLP found important the *increment* in *Flicker* compared to the baseline level for deciding that Sample 227 is a fault. Similarly, the MLP found important the *decrement* in *Minimum Power Factor* compared to the baseline level, to classify Sample 227 as a fault.

A red bar, instead, indicates that a value x_i decreases the confidence in the classifier that the sample is actually a fault,

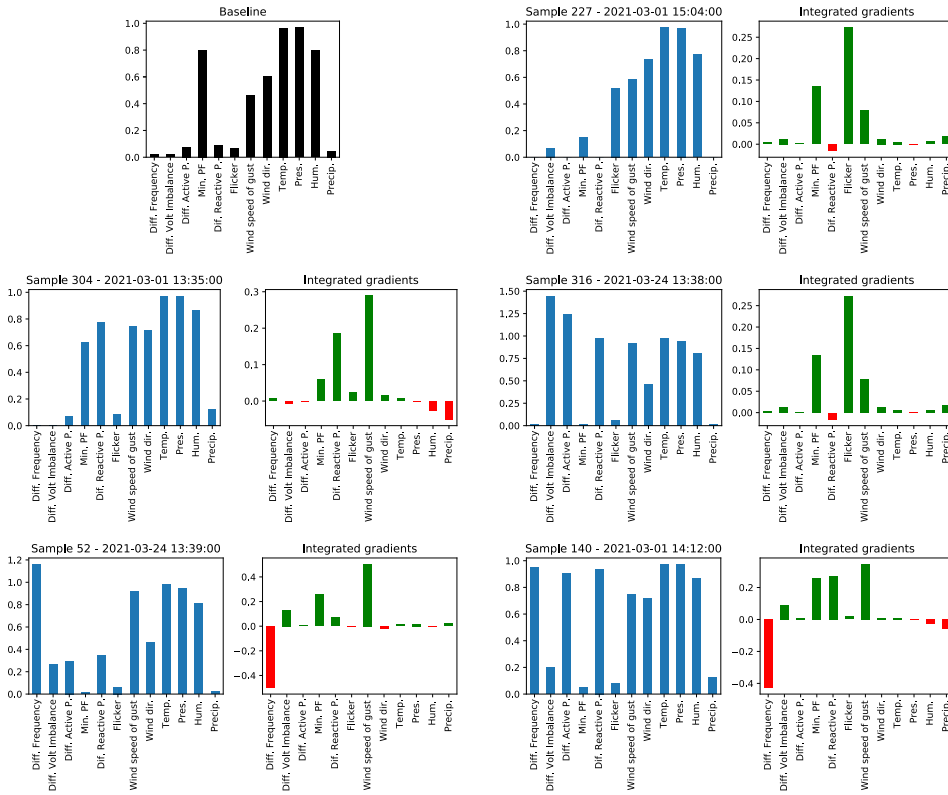


FIGURE 7. The green bars denote that a feature is important for the classification result. The higher the green bar, the more the feature value in the sample (blue bar) explains the classification result, compared to the value in the baseline (black bar). The red bars means that the value of the features in the sample decrease the confidence of the classifier that the sample is actually a fault.

TABLE 3. True positives and confidence score assigned by the MLP classifier.

	Fault ID	Confidence (MLP)	Date	Time
0	31	0.71	2021-03-02	14:30:00
1	35	0.72	2021-02-28	20:40:00
2	52	0.91	2021-03-24	13:39:00
3	140	0.92	2021-03-01	14:12:00
4	163	0.60	2021-04-07	05:16:00
5	189	0.71	2021-03-02	14:31:00
6	227	0.95	2021-03-01	15:04:00
7	235	0.88	2021-02-28	14:49:00
8	269	0.77	2021-02-28	19:28:00
9	271	0.86	2021-03-24	16:01:00
10	291	0.86	2021-03-01	14:03:00
11	304	0.92	2021-03-01	13:35:00
12	316	0.96	2021-03-24	13:38:00

compared to having a baseline value x_i' . For example, the MLP would have been even more confident that Sample 52 and Sample 140 are faults if their Difference in Frequency values would have been as in the baseline. In other words, for these two samples, the increment of Difference in Frequency is something that decreases the confidence of the classifier that they are faults.

This analysis shows that each sample has different features that are found important by the MLP for the classification.

For example, Sample 227 is classified as a fault mainly because of the above-average value in Flicker; Sample 52 is a fault due to the high value of Wind speed of gust and low value in Minimum Power Factor; for Sample 304 is important that the Difference in Reactive Power is higher than average.

The Minimum Power Factor and Reactive Power are important variables that contribute to explaining the current power quality in a power grid. The Power Factor is the ratio of the working power over the apparent power and quantifies the energy efficiency: the lower the power factor, the less efficient is the power usage of the end-customer. The Reactive Power is the amount of power dissipated in the system. A high amount of reactive power in the system could affect the power quality negatively as there will be less amount of available active power that can be used by the end-customer [57]. Therefore, it is reasonable to observe a relationship between the low value in the Minimum Power factor, and the high Difference in Reactive Power for the fault samples 52 and 304.

Interestingly, the Minimum Power Factor and Difference in Reactive Power were not emerging as important features with the global interpretability approach, which is based on the weights magnitude of the linear models. Indeed, an approach that averages the contribution of the different factors across the whole dataset is likely to conceal the importance of

configurations in the features value that appears only in a few samples. On the other hand, by analyzing samples individually, IG can reveal new patterns in the data and allow to gain deeper insights about the true causes underlying specific faults.

VI. CONCLUSION

In this work, we tackled the problem of detecting unscheduled faults in the power grid, which have major consequences for customers, such industries, relying on a stable power supply. In collaboration with the DSO, we built a data set consisting of meteorological and power data variables, which monitor potentially relevant factors to cause power faults. Once the dataset was constructed, we trained different classifiers to detect imminent faults from the value of meteorological and power variables.

The classification performance was compared in terms of F1 score and the MLP classifier achieved the top performance, followed by the Ridge Classifier. The good classification results motivated the interpretation of the decision process learned by the model, as a tool to identify the variables that mostly explain the onset of power faults. We explored two different interpretability techniques. First, we considered the magnitude of the coefficient of the linear models to quantify the importance that, on average, the different features have to determine if a sample in the dataset is a fault. The results indicated that the amount of Flicker and Wind speed of gust are the most important variables in explaining the power disturbances. Such a global interpretability approach averages the contribution of the different factors across the whole dataset and, therefore, might fail to show interesting configurations in the features value that appear only in a few samples.

As a second interpretability technique, we used the Integrated Gradients to interpret the decision process taken by the MLP classifier on individual samples. This second approach allowed us to understand what features were considered important to classify a specific sample as a fault. Interestingly, some samples were classified as faults not only for having high values in Flicker and Wind speed of gust. In fact, the IG technique showed that the MLP classified as faults samples where the Minimum Power Factor was below average or where the Difference in Reactive Power was higher than average.

The proposed interpretability techniques revealed important patterns in the data, which allow us to gain deeper insights into the underlying causes of power faults. This type of knowledge will help the DSO to give more reliable warnings to its customers (both producers and consumers) that there is an enhanced risk for grid failure when certain meteorological and power flow conditions are met. With this information, the customers can take preemptive actions to reduce the negative consequences occurring when a fault strikes.

By understanding the major factors causing the faults, the DSO will also be able to better plan how to strengthen the grid

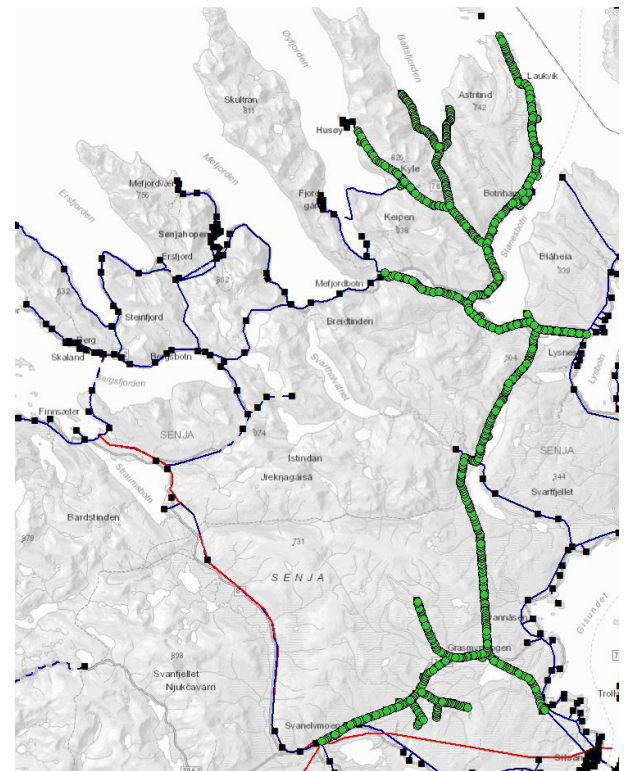


FIGURE 8. The SVAN22LY1 power grid. The power is distributed towards the north from the south. Each green dot represents a unique position of a utility pole.

to withstand incoming faults. Typical actions to improve the grid stability are: i) make changes in grid topology, such as optimizing coupling to make the grid stronger, isolating parts of the grid more likely to fail, running island mode whenever possible; ii) optimizing or even increasing the local production by introducing new power sources, including renewable ones; iii) reduce or adjust power flows by controlling flexible loads.

These kinds of strategies to mitigate incoming grid faults are currently being developed by the DSO operating the grid in our study. In particular, the local power company is installing a large battery system that should be activated right before an incoming power fault, to provide additional power and avoid instability in the power supply. Understanding which variable should be monitored to detect an incoming power fault is, therefore, fundamental to plan the installation and management of the batteries.

APPENDIX A THE INVESTIGATED POWER GRID

The power grid analyzed in this study is a radial distribution system serving an Arctic community located approximately at (69.257°N, 17.589°E). Arva Power Company, the DSO of the power grid, has named this specific grid as SVAN22LY1. Fig. 8 shows an overview of the whole SVAN22LY1 grid, indicated by green dots. The SVAN22LY1 grid spans over 60 kilometers from the south to the northernmost point and

has several branches to various communities towards the north. There are 978 unique utility poles (marked by green dots in Fig. 8) that support the power lines. The black boxes in Fig. 8 represent the electric transformer stations connected to the power grids. The red lines represent a power grid with an operating voltage of 66 kV, while the blue lines represent a power grid with an operating voltage of 22 kV. The SVAN22LY1 radial grid covered by green dots has an operating voltage of 22 kV. The largest customers connected to the SVAN22LY1 grid are located at the end of the northernmost point of the radial.

The total energy demand in the SVAN22LY1 grid is dominated by the load consumption of the local industry. The industry performs fish processing activities that are highly seasonal and uses many electrical machines in the production line that require stable power quality. Even minor power disturbances in the power supply trigger significantly long interruptions since the automated production line needs to be reset. In particular, for every short-term power interruption that occurs, is necessary to wait from 40 minutes to 1 hour before resuming the production. The consequences of the power disturbances are exacerbated by the topology of the power grid, which has a radial distribution with no alternative power supply in periods with disturbances.

APPENDIX B DATASET CONSTRUCTION

A. FAULT REPORTS

The reported faults used in this study are logged by a power-quality (PQ) metering system, which was installed in February 2021 in the proximity of the local industries to continuously measure the power quality.

The PQ meter installed by the DSO is Dranetz HDPQ-DN, and is a monitoring instrument that is used for continuous monitoring of power systems. Such PQ meters are valuable tools to get better insight and knowledge about the actual power quality. More technical details about the PQ meter are provided in [58]. The PQ metering system reports all incidents with a voltage variation of $\pm 10\%$ from the nominal values on each phase of a three-phase system with phases A, B, C. According to the standard definition, all variations of $\pm 10\%$ from normal conditions are defined as a voltage variation, and a drop larger than 10% is referred to as a voltage dip [59]. Voltage dips could provoke tripping of sensitive components such as industrial machines.

B. WEATHER MEASUREMENTS

The weather variables that are considered relevant in causing power faults are: wind speed of gust, wind direction, temperature, pressure, humidity, and precipitation. The weather data are collected from areas that are more exposed to harsh weather conditions, such as hills and cliffs near the sea coast. To collect the weather data in the Arctic region of interest, we used the AROME-Arctic weather model.[¶] This model is

developed by the meteorological institute of Norway (MET) and provides a reanalysis of historical weather data since November 2015 with a spatial resolution of 2.5 kilometers and a temporal resolution of 1 hour.

To collect the weather variables, the geographical coordinates from the weather-exposed areas in the power grid are used as inputs to the AROME-Arctic model. The output from the AROME-Arctic model is a dataset of 6 weather variables from the weather-exposed areas that are analyzed.

C. ELECTRICITY LOAD MEASUREMENTS

It is reasonable to assume that some types of fault are not caused by weather phenomena but originate from external factors that influence the power flows on the grid. To capture these effects, 6 different power-related variables from the largest industry connected to SVAN22LY1 are collected. The variables selected as relevant to explain power faults are: difference in frequency, voltage imbalance, the difference in active and reactive power, minimum power factor, and, finally, the amount of flicker in the system. All variables are metered on three different phases (phases A, B, and C).

A *change in power frequency* could be caused if there is an imbalance between energy production and consumption in the system. If there is a change in the power frequency (50 Hz is the normal frequency), the imbalance could cause power disturbances for the end-use customers.

Voltage imbalance is a voltage variation in the power system in which the voltage magnitudes or the phase angle between the different phases are not equal. It is believed that rapid changes (big changes within seconds/minutes) in power consumption at large industries could affect the power quality. Therefore, the *difference in active and reactive power* for each phase within each minute is computed. If the difference is large, there is a high activity at the industries, which are reported by the locals to result in a larger probability for faults.

The *minimum power factor* represents the relationship between the amount of active and reactive power in the system. If the minimum power factor is low, there is an increased amount of reactive power in the system. In the end, the amount of flicker in the system is collected.

Flicker is considered as a phenomenon in the power system and is closely connected to voltage fluctuations over a certain time frame [60]. A voltage fluctuation is a regular change in voltage that happens when the machinery that requires a high load is starting. In addition, rapid changes in load demand could cause voltage fluctuations. If there are several start-up situations, or the load varies significantly during a given time frame, it will be measured a high amount of flicker in the system. The amount of flicker is particularly relevant in the industry considered in this study, as they have several large machines that require high loads and have a cyclical varying load pattern. In this study, the time frame of the flicker is 10-minutes, which is the standard for measuring the short-term flicker [7].

[¶]<https://www.met.no/en/projects/The-weather-model-AROME-Arctic>

The PQ metering system has a 1-minute resolution, while the weather data have a 1-hour resolution. To align the temporal resolution of the different types of variables, the power consumption data are sub-sampled by taking one sample every 60. As an alternative sub-sampling technique, we tested taking the average of the values within each batch of 60 consecutive samples of power measurements. However, the results did not change significantly and, therefore, the former sub-sampling method was adopted.

APPENDIX C A BRIEF HISTORY OF EXPLAINABILITY IN DEEP LEARNING

Due to the presence of many non-linear transformations, it is difficult to interpret the decision process of a neural network. During the last decade, considerable research effort has been devoted towards developing insights into what a neural network learns and how it makes its decisions. While most of the explanatory techniques were originally developed in the field of computer vision, some of them can be applied also to neural networks that process sequential or vectorial data.

Gradient-based approaches aim at identifying which inputs have the most influence on the model scoring function for a given class. The pioneering work of Simonyan et al. [53] proposed to compute a saliency map by taking the gradient of the class activation score (usually, the input to the last softmax) with respect to each input feature. The visualization of the saliency maps was successively improved by using tricks such as clipping the gradients, averaging the gradients after adding Gaussian noise to the original images, and taking the absolute value of the gradients [54].

In [55], the authors propose a method to project the activations of an intermediate hidden layer back to the input space. The procedure consists in approximately inverting the operations of a CNN (affine transformations, ReLU activations, MaxPooling) from the hidden layer to the input layer. The result gives an insight into which details the hidden layer has captured from the input image.

The Guided Back Propagation approach performs the standard gradient back-propagation but, when a ReLU is encountered, the gradient is back-propagated only if both the gradient and the ReLU activation in the forward pass are positive [56].

As a drawback, gradient-based methods attribute zero contribution to inputs that saturate the ReLU or MaxPool. To capture such shortcomings, a formal notion of explainability (or relevance) was introduced in [51]. In particular, the authors introduced an axiom on the *conservation of total relevance*, which states that the sum of the relevance of all pixels must equal the class score of the model. The authors propose to distribute the total relevance of the class score to the input features with a method called Layer-wise Relevance Propagation (LRP). The class score is computed as the difference between the score obtained by the actual input and the score obtained by an uninformative input, called *baseline*. Each time the relevance is passed down from a neuron to the


contributing neurons in the layer below, the total relevance of contributing neurons is preserved. All incoming relevances to a neuron from the layer above are collected and summed up before passing down further. By doing this recursively from layer to layer, the input layer is eventually reached, which gives the relevance of each input feature. The relevance of a neuron to its contributing inputs can be distributed based on the magnitude of the weights of the neural network layers.

While LRP followed the conservation axiom, it did not formalize how to distribute the relevance among the input features. To address this problem DeepLiFT [52] enforces an additional axiom on how to propagate the relevance down, by following the chain rule like gradients.

REFERENCES

- [1] S. Chiaradonna, F. Di Giandomenico, and G. Masetti, "Analyzing the impact of failures in the electric power distribution grid," in *Proc. 7th Latin-American Symp. Dependable Comput. (LADC)*, Oct. 2016, pp. 99–108.
- [2] C. Klinger, O. Landeg, and V. Murray, "Power outages, extreme events and health: A systematic review of the literature from 2011–2012," *PLoS Currents*, to be published.
- [3] T. H. Meles, "Impact of power outages on households in developing countries: Evidence from Ethiopia," *Energy Econ.*, vol. 91, Sep. 2020, Art. no. 104882.
- [4] M. Shuai, W. Chengzhi, Y. Shiwen, G. Hao, Y. Jufang, and H. Hui, "Review on economic loss assessment of power outages," *Proc. Comput. Sci.*, vol. 130, pp. 1158–1163, May 2018.
- [5] S. Tully, "The human right to access electricity," *Electr. J.*, vol. 19, no. 3, pp. 30–39, Apr. 2006.
- [6] G. S. Gopinath and M. Meher, "Electricity a basic need for the human beings," *AIP Conf. Proc.*, vol. 1992, no. 1, 2018, Art. no. 040024.
- [7] N. Lovdata. *Forskrift Om Leveringskvalitet I Kraftsystemet*. [Online]. Available: <https://lovdata.no/dokument/SF/forskrift/2004-11-30-1557>
- [8] C. Rubí. *The Challenges of Upgrading the Power Grid for a Decarbonised Electric Future*. [Online]. Available: <https://informaconnect.com/the-challenges-of-upgrading-the-power-grid-f%or-a-decarbonised-electric-future/>
- [9] N. Sapountzoglou, J. Lago, B. De Schutter, and B. Raison, "A generalizable and sensor-independent deep learning method for fault detection and location in low-voltage distribution grids," *Appl. Energy*, vol. 276, Oct. 2020, Art. no. 115299.
- [10] K. Chen, J. Hu, Y. Zhang, Z. Yu, and J. He, "Fault location in power distribution systems via deep graph convolutional networks," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 1, pp. 119–131, Jan. 2020.
- [11] H. Khorasgani, A. Hasanzadeh, A. Farahat, and C. Gupta, "Fault detection and isolation in industrial networks using graph convolutional neural networks," in *Proc. IEEE Int. Conf. Prognostics Health Manage. (ICPHM)*, Jun. 2019, pp. 1–7.
- [12] R. M. G. Ferrari, T. Parisini, and M. M. Polycarpou, "Distributed fault detection and isolation of large-scale discrete-time nonlinear systems: An adaptive approximation approach," *IEEE Trans. Autom. Control*, vol. 57, no. 2, pp. 275–290, Feb. 2012.
- [13] E. Balouji, I. Y. H. Gu, M. H. J. Bollen, A. Bagheri, and M. Nazari, "A LSTM-based deep learning method with application to voltage dip classification," in *Proc. 18th Int. Conf. Harmon. Quality Power (ICHQP)*, May 2018, pp. 1–5.
- [14] F. M. Bianchi, E. De Santis, A. Rizzi, and A. Sadeghian, "Short-term electric load forecasting using echo state networks and PCA decomposition," *IEEE Access*, vol. 3, pp. 1931–1943, 2015.
- [15] F. M. Bianchi, E. Maiorino, M. C. Kampffmeyer, A. Rizzi, and R. Jenssen, "Recurrent neural networks for short-term load forecasting: An overview and comparative analysis," *Tech. Rep.*, 2017.
- [16] A. T. D. Perera, V. M. Nik, D. Chen, J.-L. Scartezini, and T. Hong, "Quantifying the impacts of climate change and extreme climate events on energy systems," *Nature Energy*, vol. 5, no. 2, pp. 150–159, Feb. 2020.
- [17] H. Sabouhi, A. Doroudi, M. Fotuhi-Firuzabad, and M. Bashiri, "Electrical power system resilience assessment: A comprehensive approach," *IEEE Syst. J.*, vol. 14, no. 2, pp. 2643–2652, Jun. 2020.

- [18] D. N. Trakas, M. Panteli, N. D. Hatziaargyriou, and P. Mancarella, "Spatial risk analysis of power systems resilience during extreme events," *Risk Anal.*, vol. 39, no. 1, pp. 195–211, Jan. 2019.
- [19] M. Panteli and P. Mancarella, "Influence of extreme weather and climate change on the resilience of power systems: Impacts and possible mitigation strategies," *Electr. Power Syst. Res.*, vol. 127, pp. 259–270, Oct. 2015.
- [20] F. de Caro, E. M. Carlini, and D. Villacci, "Flexibility sources for enhancing the resilience of a power grid in presence of severe weather conditions," in *Proc. AEIT Int. Annu. Conf. (AEIT)*, Sep. 2019, pp. 1–6.
- [21] D. Owerko, F. Gama, and A. Ribeiro, "Predicting power outages using graph neural networks," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Nov. 2018, pp. 743–747.
- [22] O. F. Eikeland, F. M. Bianchi, I. S. Holmstrand, S. Bakkejord, S. Santos, and M. Chiesa, "Uncovering contributing factors to interruptions in the power grid: An Arctic case," Tech. Rep., 2021.
- [23] O. P. Mahela, A. G. Shaik, and N. Gupta, "A critical review of detection and classification of power quality events," *Renew. Sustain. Energy Rev.*, vol. 41, pp. 495–505, Jan. 2015.
- [24] M. Mishra, "Power quality disturbance detection and classification using signal processing and soft computing techniques: A comprehensive review," *Int. Trans. Electr. Energy Syst.*, vol. 29, no. 8, Aug. 2019, Art. no. e12008.
- [25] K. Michalowska, V. Hoffmann, and C. Andresen, "Impact of seasonal weather on forecasting of power quality disturbances in distribution grids," in *Proc. Int. Conf. Smart Energy Syst. Technol. (SEST)*, Sep. 2020, pp. 1–6.
- [26] F. Yang, D. W. Wanik, D. Cerrai, M. A. E. Bhuiyan, and E. N. Anagnostou, "Quantifying uncertainty in machine learning-based power outage prediction model training: A tool for sustainable storm restoration," *Sustainability*, vol. 12, no. 4, p. 1525, Feb. 2020.
- [27] A. M. Salman and Y. Li, "A probabilistic framework for multi-hazard risk mitigation for electric power transmission systems subjected to seismic and hurricane hazards," *Struct. Infrastruct. Eng.*, vol. 14, no. 11, pp. 1499–1519, Nov. 2018.
- [28] A. M. Salman and Y. Li, "A probabilistic framework for seismic risk assessment of electric power systems," *Proc. Eng.*, vol. 199, pp. 1187–1192, Jan. 2017.
- [29] A. M. Salman and Y. Li, "Multihazard risk assessment of electric power systems," *J. Struct. Eng.*, vol. 143, no. 3, Mar. 2017, Art. no. 04016198.
- [30] S. Mukherjee, R. Nateghi, and M. Hastak, "A multi-hazard approach to assess severe weather-induced major power outage risks in the U.S.," *Rel. Eng. Syst. Saf.*, vol. 175, pp. 283–305, Jul. 2018.
- [31] R. Eskandarpour and A. Khodaei, "Machine learning based power grid outage prediction in response to extreme events," *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 3315–3316, Jul. 2017.
- [32] K. Murray and K. R. W. Bell, "Wind related faults on the GB transmission network," in *Proc. Int. Conf. Probabilistic Methods Appl. Power Syst. (PMAPS)*, Jul. 2014, pp. 1–6.
- [33] V. Hoffmann, K. Michalowska, C. Andresen, and B. N. Torsæter, "Incipient fault prediction in power quality monitoring," Tech. Rep., 2019.
- [34] K. W. Hoiem, V. Santi, B. N. Torsater, H. Langseth, C. A. Andresen, and G. H. Rosenlund, "Comparative study of event prediction in power grids using supervised machine learning methods," in *Proc. Int. Conf. Smart Energy Syst. Technol. (SEST)*, Sep. 2020, pp. 1–6.
- [35] G. M. Abusdal, G. T. Heydt, and A. Ripegut, "Utilization of advanced metering infrastructure in back-fed ground fault detection," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Jul. 2015, pp. 1–5.
- [36] T. S. Tyvold, B. Nybakk Torsater, C. A. Andresen, and V. Hoffmann, "Impact of the temporal distribution of faults on prediction of voltage anomalies in the power grid," in *Proc. Int. Conf. Smart Energy Syst. Technol. (SEST)*, Sep. 2020, pp. 1–6.
- [37] G. H. Rosenlund, K. W. Hoiem, B. N. Torsater, and C. A. Andresen, "Clustering and dimensionality-reduction techniques applied on power quality measurement data," in *Proc. Int. Conf. Smart Energy Syst. Technol. (SEST)*, Sep. 2020, pp. 1–6.
- [38] R. Fainti, M. Alamaniotis, and L. H. Tsoukalas, "Three-phase line overloading predictive monitoring utilizing neural networks," in *Proc. 19th Int. Conf. Intell. Syst. Appl. Power Syst. (ISAP)*, Sep. 2017, pp. 1–6.
- [39] Y. Zhou, R. Arghandeh, and C. J. Spanos, "Partial knowledge data-driven event detection for power distribution networks," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 5152–5162, Aug. 2018.
- [40] K. Manivinnan, C. L. Benner, B. D. Russell, and J. A. Wischkaemper, "Automatic identification, clustering and reporting of recurrent faults in electric distribution feeders," in *Proc. 19th Int. Conf. Intell. Syst. Appl. Power Syst. (ISAP)*, Sep. 2017, pp. 1–6.
- [41] S. Khokhar, A. A. M. Zin, A. P. Memon, and A. S. Mokhtar, "A new optimal feature selection algorithm for classification of power quality disturbances using discrete wavelet transform and probabilistic neural network," *Measurement*, vol. 95, pp. 246–259, Jan. 2017.
- [42] O. Zybalkina, M. Domagk, J. Meyer, and P. Schegner, "A feature-based method for automatic anomaly identification in power quality measurements," in *Proc. IEEE Int. Conf. Probabilistic Methods Appl. Power Syst. (PMAPS)*, Jun. 2018, pp. 1–6.
- [43] R. A. Shuvro, P. Das, M. M. Hayat, and M. Talukder, "Predicting cascading failures in power grids using machine learning algorithms," in *Proc. North Amer. Power Symp. (NAPS)*, Oct. 2019, pp. 1–6.
- [44] C. M. Bishop, "Pattern recognition," *Mach. Learn.*, vol. 128, no. 9, 2006.
- [45] B. E. Boser, I. M. Guyon, and V. N. Vapnik, "A training algorithm for optimal margin classifiers," in *Proc. 5th Annu. Workshop Comput. Learn. Theory (COLT)*, 1992, pp. 144–152.
- [46] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [47] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [48] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [49] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent.*, 2014, pp. 1–15.
- [50] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3319–3328.
- [51] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PLoS ONE*, vol. 10, no. 7, Jul. 2015, Art. no. e0130140.
- [52] A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3145–3153.
- [53] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *Proc. Int. Conf. Learn. Represent.*, 2014, pp. 1–8.
- [54] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, "Smoothgrad: Removing noise by adding noise," in *Proc. Int. Conf. Mach. Learn.*, 2017.
- [55] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Eur. Conf. Comput. Vis.*, Springer, 2014, pp. 818–833.
- [56] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," in *Proc. Int. Conf. Learn. Represent.*, 2015, pp. 1–14.
- [57] A. Mansour, Z. Chengning, and H. Nasry, "Measurement of power components in balanced and unbalanced three-phase systems under nonsinusoidal operating conditions by using IEEE standard 1459–2010 and Fourier analysis," in *Proc. Int. Conf. Technological Adv. Electr., Electron. Comput. Eng. (TAECE)*, May 2013, pp. 166–171.
- [58] D. Technologies. *Permanent PQ and Energy Monitoring Products*. [Online]. Available: <https://www.dragnetz.com/product/hdpq-dn/>
- [59] E. Csanyi. *Detailed Overview of Power System Disturbances (Causes and Impacts)*. [Online]. Available: <https://electrical-engineering-portal.com/detailed-overview-of-power-system-disturbances-causes-and-impacts>
- [60] A. C. S. Mark Halpin, *Power Electronics Handbook*, 2nd ed. Montgomery, AL, USA: Academic, Auburn Univ., Department of Electrical and Computer Engineering, 2007, ch. 38.



ODIN FOLDVIK EIKELAND received the M.S. degree in energy, climate, and environment from the University of Tromsø, in 2019, where he is currently pursuing the Ph.D. degree. He is also a Civil Engineer in renewable energy, and took a Ph.D. related to renewable energy and energy analytics. His main research interests include renewable energy with a special focus on solar energy, time-series forecasting, neural networks, and power grid fault detection with machine learning techniques.



INGA SETSÅ HOLMSTRAND received the M.Sc. degree from the Department of Physics and Technology, University of Tromsø, in 2018. She started at Arva Power Company, in 2020, where she currently works with data analysis of the power grid.



MATTEO CHIESA is currently a Material Scientist working on creating and implementing technologies necessary to adapt the current energy system into a competitive and secure one. He design and investigate engineered nanomaterials for solar energy conversion and storage systems. His main research interest includes tackling specific challenges related to the deployment of renewable energy under arctic conditions through synthetic and symbiotic approach.



SIGURD BAKKEJORD received the M.Sc. degree in electrical power engineering from the Norwegian University of Science and Technology, in June 2011. He is currently working as a Power Grid System Manager of the Distribution System Operator Arva AS. His main research interests include power grid analysis and grid-information systems.



FILIPPO MARIA BIANCHI received the Ph.D. degree from the Department of Information Engineering, Electronics, and Telecommunications, Sapienza University of Rome, in 2016. He is currently an Associate Professor with the Department of Mathematics and Statistics, UiT-The Arctic University of Norway, and also a Research Scientist at NORCE. His research interests include machine learning, complex networks, and dynamical systems.

...