

Article

Interpretable Detection of Partial Discharge in Power Lines with Deep Learning

Gabriel Michau , Chi-Ching Hsu  and Olga Fink ^{*,†} 

Swiss Federal Institute of Technology, ETH Zürich, 8093 Zürich, Switzerland; gmichau@ethz.ch (G.M.); hsuch@student.ethz.ch (C.-C.H.)

* Correspondence: ofink@ethz.ch

† Current address: Chair of Intelligent Maintenance Systems, HIL H 27.1, Stefano-Franscini-Platz 5, 8093 Zürich, Switzerland.

Abstract: Partial discharge (PD) is a common indication of faults in power systems, such as generators and cables. These PDs can eventually result in costly repairs and substantial power outages. PD detection traditionally relies on hand-crafted features and domain expertise to identify very specific pulses in the electrical current, and the performance declines in the presence of noise or of superposed pulses. In this paper, we propose a novel end-to-end framework based on convolutional neural networks. The framework has two contributions: First, it does not require any feature extraction and enables robust PD detection. Second, we devise the pulse activation map. It provides interpretability of the results for the domain experts with the identification of the pulses that led to the detection of the PDs. The performance is evaluated on a public dataset for the detection of damaged power lines. An ablation study demonstrates the benefits of each part of the proposed framework.

Keywords: partial discharges; power distribution lines; temporal CNN; fault detection



Citation: Michau, G.; Hsu, C.-C.; Fink, O. Interpretable Detection of Partial Discharge in Power Lines with Deep Learning. *Sensors* **2021**, *21*, 2154. <https://doi.org/10.3390/s21062154>

Academic Editors: Fabio Viola and Biswanath Samanta

Received: 15 February 2021

Accepted: 17 March 2021

Published: 19 March 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Many critical services in our society, such as healthcare, transportation, and security, require a robust, reliable, and uninterrupted supply of electricity. This requires on the one hand a reliable and redundant infrastructure, and on the other hand, the ability to maintain the performance of the infrastructure. As such, the ability to detect faults and to act accordingly is of critical importance for maintaining a high availability of the network [1]. Many components of the power generation and distribution network can be directly monitored with specific sensors. However, it is not possible nor cost efficient for all the components and all the fault types. Therefore, for some of the components, the monitoring can only be performed indirectly through the behavior of the electrical current. For example, insulation damages in power systems, such as generators or defects in medium-voltage cables [2], can be monitored by detecting localized pulses in the electrical current, namely, partial discharges (PD).

According to the IEC 60270 international standard, partial discharges are “localized electrical discharges that only partially bridge the insulation between conductors”. In fact, the presence of PD can be indicative of anomalies in many electrical systems and can cause further degradation of the insulation. The high-voltage discharges deteriorate the insulation materials and can have impacts on the entire system. Their detection is, therefore, of utmost importance to assess the condition of electrical components and has been a long-standing challenge [3]. As such, the literature is extremely vast. PD detection has been studied in many systems such as in transformers [4], gas-insulated high-voltage switchgear [5], power plants [6], and power lines [7]. The main challenge in PD detection lies in the detection of extremely short and temporally localized events: their wavelength is at the micro-second scale. It requires, therefore, extremely high-frequency data (several tens of MHz). In addition, only few pulses can occur per period of the current utility

frequency (usually 50 or 60 Hz) [3]. In brief, PD signs in the electrical current represent roughly 1/20,000-th of the data. Until the very recent development of technologies able to capture and store such vast amount of data, the detection of PD patterns had to be performed online.

Among the traditional approaches, a group of approaches takes advantage of the property that for some systems, PD always occurs at the same phase in the electrical current. These approaches are also referred to as the phase-resolved partial discharge (PRPD) detection methods [8]. They consist in detecting pulses in the electrical current, whereby the simplest method to detect a pulse is to apply a maximum filter. Subsequently, the detected rate of occurrence (n) of pulses is plotted as a function of their voltage amplitude (Q) and of their phase value (ϕ). It can be implemented online and experts can inspect PRPD graphs to recognize the patterns generated by the different types of PD [9]. However, to obtain meaningful occurrence rates, these methods require the aggregation of pulses over several hundreds or thousands of periods. Overall, these methods are expensive, as experts need to constantly monitor the ϕ - Q - n diagrams. The interpretation of the diagrams also becomes difficult in the presence of noise or of superimposed pulses.

As all types of PD are not necessarily resolved in the phase domain, statistical approaches are also frequently used [10]. They aim at characterizing the pulses with engineered features. The features become a multidimensional space where decision boundaries are established, for example with traditional classifiers such as random forest [7], support vector machines (SVM) [11,12], or deep learning methods such as artificial neural networks (ANN) [13], convolutional neural networks (CNN) [14,15], autoencoders [16], and recurrent neural networks (long short-term memory networks (LSTM)) [17–19]. For a more exhaustive overview, the reader is referred to the literature reviews in [20,21]. These methods are time consuming and expensive due to the difficulty of extracting and engineering the relevant features. It requires years of domain expertise for a profound understanding of the system. Additionally, the methods suffer from degradation performances when pulses are superimposed.

To address the aforementioned challenges, we propose to take advantage of the recent advances in Deep Learning (DL) applied to time series (TS) anomaly detection. For example, DL has already been successfully applied to identify cardiac abnormalities in electrocardiography (ECG) data [22]. In particular, convolutional neural networks (CNN) have recently demonstrated very good performances for TS classification [23,24], forecasting [25], and anomaly detection [26]. In fact, temporal convolutions are able to learn meaningful filters in the time domain, adjusted to the signature of the analyzed events, and are, thus, often compared to learnable spectral features [27].

In this paper, we propose an end-to-end learning framework for partial discharge detection in time series. The framework comprises two parts: (1) the automatic PD detection without any feature engineering and (2) the subsequent extraction of pulse activation maps that provide the domain experts a possibility to interpret the results. The difficulty met by previous research in the detection of partial discharges lies in the discrimination between PD and non-PD related pulses. Therefore, we propose here to extract a collection of pulses for each period of the utility frequency. We train a temporal convolution neural network with the binary information of whether PD are present in the original time series or not. Since all pulses of the collection are processed with the same temporal filters, a well-performing model should be able learn the PD pulse signatures. Learning these signatures is, in fact, particularly valuable for the experts to potentially distinguish between different types of PD. Therefore, we design our framework to provide both competitive results in terms of partial discharge detection, and a visualization of the neural network processing of the inputs through the pulse activation maps. These activation maps provide interpretability and explainability of the results and allow the experts to diagnose for each time series, which pulses and which part of the pulses were dominant in the final score of the network. It gives an extremely fine interpretation of the network decision. We demonstrate the performance of the proposed approach by achieving rank 4 on private

leaderboard of the Kaggle VSB Power Line Fault Detection competition. The aim is to identify damaged power lines from the observed PD in the electrical voltage [2,28]. Furthermore, we also demonstrate the added value of each part of the proposed framework with an ablation study.

The paper is organized as follows. Section 2 presents each step of the framework, including the preprocessing, network architecture, and pulse activation map devising. Section 3 details the experiments we performed to quantify the results presented in Section 4. The final results are discussed in Section 5.

2. Methodology

An overview of the proposed framework is presented in Figure 1. For each measurement (1), each phase is handled independently. Low frequencies, in particular the utility frequency, are filtered first (2). Pulses are identified and ranked with a simple maximum filter and extracted into a pulse collection (3). Each phase is estimated independently by the same neural network (4). The final decision on the power line takes the results from the three phases into account and applies a global threshold (5). Last, a pulse activation map (6) is computed to understand which part of the input led to the networks classification results.

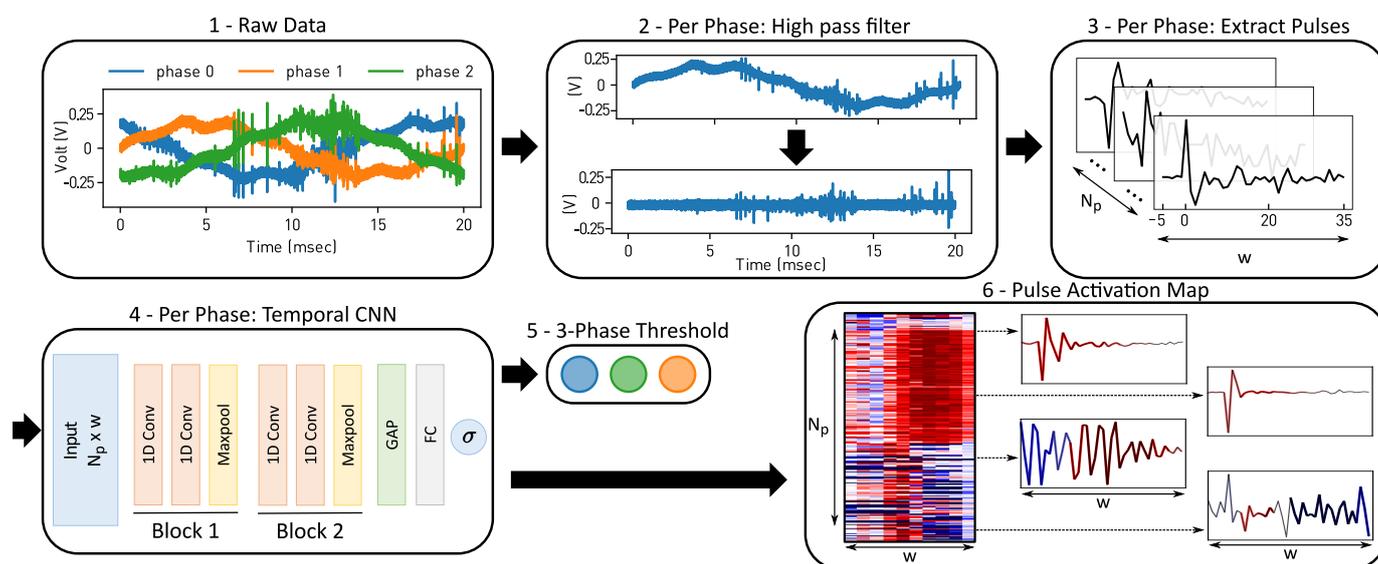


Figure 1. Graphical Abstract: (1) The raw data: three phases over one period. (2) For each phase, the signal is filtered to remove the utility frequency f_{ut} . (3) Extracted w samples of pulses identified with a maximum filter. The windows start $1/8$ th before the identified peak. (4) For each phase, the pulse collection is fed to a deep 1D convolutional neural network. (5) The outputs of the three phases are considered together to label the powerline. (6) Thanks to the Global Average Pooling (GAP) Layer, the Pulse Activation Map can be displayed for further diagnostics.

2.1. Preprocessing

The main assumption of the proposed framework relies on the inherent definition of a PD signature as a pulse in the electric current. Thus, inspired by the PRPD analysis, in the very first step, we identify and extract the pulses with a simple maximum filter. This requires first the removal of the low frequencies.

Data filtering—Figure 1 (2). PDs are due to insulation failures and typically occur at very specific voltage changes. Their frequency is much higher than the utility frequency f_{ut} . Thus, we first implement a high-pass filter with cut-off frequency $f_c > f_{ut}$. In this work, we apply a Butterworth filter of order 5 [29]. However, other filters could be explored. An example before and after the filtering of an electric current recording over one period is shown in Figure 1 (2), where the sampling frequency $f_s = 40$ MHz, $f_c = 20$ kHz and $f_{ut} = 50$ Hz. Low frequencies such as the underlying sine wave are eliminated after the high-pass filter, while the high frequency pulses remain unchanged.

Pulse extraction—Figure 1 (3). As the partial discharge signature is inherently a pulse in the electric potential, we propose as a second step to extract a large collection of pulses from the recordings that will be used as inputs to the neural network (NN). The goal of the NN is to learn to recognize if there is a PD pulse signature within the input collection. Due to the nature of partial discharge, we can expect some periodicity in the occurrence of the pulses with respect to the utility frequency. We create therefore, for each period of the utility frequency, a 2D array where each row represents a single pulse. The columns represent the time dimension. The number of columns corresponds to the number of timestamps w collected for each pulse.

The pulses are identified with a maximum filter on the absolute value of the electric potential. The filter extracts the local maximal values which are further apart than a given window size. For simplicity purposes, we set this window to w . We extract in the filtered data, the w timestamps around each of the N_p largest local maxima, with an offset of $1/8 \cdot w$. That is, if the i -th local maximum is localized at timestamp t_i , we extract the interval $[t_i - \frac{w}{8}, t_i + \frac{7w}{8}]$. The collection of pulses is, therefore, a 2D array of shape $N_p \times w$. Figure 1 (3) illustrates the pulse extraction and the resulting collection for one period and phase of the utility frequency.

N_p and w are hyperparameters of the proposed approach. Expert domain knowledge could help to identify relevant values for these hyperparameters. The selected values would primarily depend on the noise level for the selection of N_p (some noise pulses may dominate the PD pulses), and on the expected frequency of the pulses for the selection of w . Yet, if this knowledge is not available, as in our case, these hyperparameters can be selected based on cross-validation.

2.2. Temporal Convolutional Neural Network

Convolutional neural network (CNN). Inspired by the recent advances in computer vision and more recently also in time series classification tasks, using CNN architectures, we propose to apply a CNN for this PD detection task. For the applied architecture, we propose a deep learning neural network architecture with a similar structure to VGG neural networks [30]. Yet, the architecture requires several adaptations to the high-frequency time series data as used in this case study. Unlike in images, where the neighborhood of a pixel has a clear meaning in both the X and the Y dimensions, in the extracted pulses, only the temporal dimension contains a physically meaningful neighborhood relationship. The pulses have been ordered by decreasing amplitudes and not by their relationships in the signal. We, therefore, apply 1D convolutions instead of 2D kernels. This also means that the temporal filters are applied identically to each pulse, performing operations similar to spectral analysis.

Global Average Pooling (GAP). A limitation when using CNN is that the convolutional layers preserve the dimensionality of their inputs. Therefore, as predictions are usually vectors (with one element per class), it is necessary to flatten the latent space in order to transition toward fully connected layers. A consequence is an explosion in the number of model parameters, often leading to overfitting effects and harming the generalization ability of the network. We propose, therefore, to use the Global Average Pooling (GAP) as a structural regularizer [31,32]. GAP takes the average over the feature maps channel-wise and thus shrinks the size of the last latent space before its vectorization.

Proposed CNN Architecture. The proposed network architecture takes advantage of 1D CNN and the GAP layer. It contains 2 blocks comprising 2 successive convolutional layers and a max pooling layer. The 2 blocks are followed by a GAP layer, a fully connected (FC) layer, and a single neuron layer for binary classification.

2.3. Pulse Activation Maps (PAM)

To provide more interpretability and more insights on the network decision to classify a collection of pulses as belonging to a damaged line or not, we propose to devise the Class Activation Maps (CAM) of our network [33]. Following the methodology in [33], we devise

in this section the pulse activation map (PAM) for the proposed network architecture. The PAM enables to interpret which part of the pulse has contributed most to the classification result (in this case, PD, or non-PD). There are two differences to the original contribution. First, our network has a binary output. We devise a single PAM per input, instead of a per-class activation map. Second, our network contains a fully connected layer between the GAP and the output. In the following, we demonstrate that the CAM (or here, the PAM) can still be computed in such cases, as long as the activation functions used by the intermediate fully connected layers are piece-wise linear.

As the first two blocks of our network use 1-D convolution filters, the latent space M after the last block can be written as $M = \left(m_{plj} \right)_{p \in \llbracket 1, N_p \rrbracket, l \in \llbracket 1, N_l \rrbracket, j \in \llbracket 1, N_f \rrbracket}$, where N_p is the number of pulses used as input to the network, N_l is the resulting size after the successive max-pooling (here $N_l = \lfloor \frac{w}{4} \rfloor$), and N_f is the number of filters of the last convolution. In the following, we denote the size of M as $N_m = N_p \cdot N_l$

The j -th neuron of the GAP layer performs the operation:

$$\forall j \in \llbracket 1, N_f \rrbracket, \quad \text{GAP}^j = \frac{1}{N_m} \sum_{p,l}^{N_p \times N_l} m_{plj} \quad (1)$$

The GAP layer is connected to a fully connected layer of size N_{FC} with the weights and bias $(w_{ij}^{FC}, b_i^{FC})_{i \in \llbracket 1, N_{FC} \rrbracket, j \in \llbracket 1, N_f \rrbracket}$, we have for the i -th neuron, before activation:

$$\begin{aligned} \forall i \in \llbracket 1, N_{FC} \rrbracket, \\ \text{FC}^i &= b_i^{FC} + \sum_{j=1}^{N_f} w_{ij}^{FC} \cdot \frac{1}{N_m} \sum_{p,l}^{N_p \times N_l} m_{plj} \\ &= \frac{1}{N_m} \sum_{p,l}^{N_p \times N_l} \left[b_i^{FC} + \sum_{j=1}^{N_f} w_{ij}^{FC} \cdot m_{plj} \right] \end{aligned} \quad (2)$$

After the piece-wise linear activation function, the i -th activated neuron AFC^i is given by

$$\begin{aligned} \text{AFC}^i &= \text{ReLU}(\text{FC}^i) \\ &= \delta_i \cdot \text{FC}^i \end{aligned} \quad (3)$$

where

$$\delta_i = \begin{cases} 1 & \text{if } \text{FC}^i > 0 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Note that the definition of δ can be generalized to any piece-wise linear activation function. Under this assumption, it is also worth noticing that Equation (3) can be generalized to as many successive fully-connected layers as required.

Last, the activated neurons are combined into a last dense layer of output size 1. Denoting its weights and bias by $(w_i^{Out}, b^{Out})_{i \in \llbracket 1, N_{FC} \rrbracket}$. The final score before the sigmoid is obtained as

$$\begin{aligned} \text{Score} &= b^{Out} + \sum_{i=1}^{N_{FC}} w_i^{Out} \cdot \text{AFC}^i \\ &= \frac{1}{N_m} \sum_{p,l}^{N_p \times N_l} \left[b^{Out} + \sum_{i=1}^{N_{FC}} w_i^{Out} \delta_i \left(b_i^{FC} + \sum_{j=1}^{N_f} w_{ij}^{FC} m_{plj} \right) \right] \end{aligned} \quad (5)$$

Finally, the pulse activation map PAM for each input is a collection of N_p vectors of size N_l and is defined as

$$\forall (p, l) \in \llbracket 1, N_p \rrbracket \times \llbracket 1, N_l \rrbracket, \quad \text{PAM}(p, l) = b^{\text{Out}} + \sum_{i=1}^{N_{\text{FC}}} w_i^{\text{Out}} \delta_i \left(\mathbf{b}_i^{\text{FC}} + \sum_{j=1}^{N_f} w_{ij}^{\text{FC}} m_{plj} \right) \quad (6)$$

As the decision of the network is taken after applying the sigmoid operation to the Score value Equation (5), we can interpret the PAM as follows. A map whose average is positive corresponds to a score above 0.5 after the sigmoid operation, and is thus originating from pulses containing PD. On the contrary, a map whose average is negative corresponds to a non-PD pulse. The activation maps can be used by domain experts to further evaluate the pulses and possibly to distinguish between different types of PD.

3. Experiments

3.1. Datasets

To demonstrate the benefit of our approach, we apply the proposed methodology on the VSB dataset, generated and released by the Technical University of Ostrava [28]. The goal of the case study is to detect damaged three-phase, medium-voltage overhead power lines [2]. According to the dataset description, damaged power lines can be identified through the observed PD patterns [28]. To this end, the electric voltage is recorded over one period of the grid utility frequency, 50 Hz, for the three phases simultaneously. The sampling frequency is 40 MHz such that each recording contains 800,000 values. An example signal is shown in Figure 1 (1).

The VSB dataset contains two sets of measurements. The training set contains 8712 samples with 3 labels: the measurement ID, the phase, and whether the power line insulation was damaged at the time of recording. Damaged power lines should contain PD, however no additional information is provided on the PD types, shapes, or location. In this set, 575 samples are labeled as damaged power lines.

The second set contains 20,037 samples with two labels: the measurement ID and the phase. No ground truth is provided with respect to the presence of PD. However, the predictions of the health state can be evaluated online through the Matthew Correlation Coefficient (MCC).

To the best of our knowledge, no other published study outside of the competition leaderboard reported results on the second test dataset. In [12,18,19], the reported results are computed on a subset of the labeled dataset. In [12], results are reported on the full training set and might therefore be overfitted. In [18], results are reported on an artificially augmented set containing 807 non-PD signals and 935 signals with PD, which might also therefore suffer overfitting. We report anyway their results in Table 2, where we recompute the value of the metrics they would achieve on our set, assuming constant sensitivity and specificity of their model. This cannot be done for the work in [19], as the numbers of tested samples with and without PD are not reported.

3.2. Network Architecture and Training

The proposed neural network architecture as presented in Figure 1 (4) comprises two convolutional blocks: The first block contains two temporal convolutional layers with 16 kernels of size 15. The second block has two temporal convolutional layers with 8 kernels of size 10. Each block is followed by a 1D temporal max-pooling layer with kernel size 2. Therefore, the input size is $N_p \times w \times 1$, the latent space size after the first block is $N_p \times \frac{w}{2} \times 16$, the latent space size after the second block is $N_p \times \frac{w}{4} \times 8$, and the latent space size after the GAP layer is 8. In particular, we have $N_l = \frac{w}{4} = 10$ in Equation (6). The fully connected layer after the GAP layer is of size 32. All layers but the last output layer use *ReLU* as activation function. The hyperparameters of this architecture (number of blocks, kernel number, and size) were inferred from a grid search with a 5-fold stratified cross-validation.

We implemented the network with Keras and TensorFlow. For the training, we used the ADAM [34] optimizer with constant learning rate of 10^{-3} , $\beta_1 = 0.9$, and $\beta_2 = 0.999$. We used the binary cross entropy loss:

$$L = -(y \log(p) + (1 - y) \log(1 - p)), \quad (7)$$

where y is the ground truth and p is the network output.

3.3. Threshold Setting

The problem at hand is a binary classification problem. The output is, therefore, designed as a single neuron output layer, activated with a sigmoid function such that the output is continuous between 0 and 1. The traditional baseline consists in using a threshold at value 0.5 on the network output value. Compared to the baseline approach, we propose to explore two modifications: first, the inference of an optimized threshold th_v based on a validation set, and second, the consideration of the three phases as a single indicator of the power line health. We propose, thus, to compare four different postprocessing approaches of the network output:

- (i) Baseline: Round the output (threshold $th = 0.5$).
- (ii) '1=3'-Phase Classification: Round the prediction of the three phases. If one phase is estimated as damaged, the whole power line (the three phases) is considered as being damaged.
- (iii) 1-Phase Optimized Threshold: Infer the threshold th_v with cross-validation, and consider each phase independently.
- (iv) Proposed 3-Phase Global threshold: Using the threshold th_v devised in (iii), apply the 3-Phase Global threshold, $G_{th_v} = 3 \cdot th_v$, to the sum of the output values for the three phases.

Please note that in (iv), the direct inference of the G_{th_v} with cross-validation instead of using $3 \cdot th_v$ may have enhanced the performance. This is, however, not possible as in the training set, some samples do not have all the three phases recorded as damaged. Using 5-fold stratified cross validation to maximize the MCC (defined below in Section 3.4), we define $th_v = 0.28$.

3.4. Evaluation Metrics

As part of the competition, a tool is provided to evaluate the test set online. It provides an evaluation of the results with the Matthews Correlation Coefficient (MCC) [35]:

$$MCC = \frac{(TP \cdot TN) - (FP \cdot FN)}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (8)$$

where TP is the number of true positives, TN is the number of true negatives, FP is the number of false positives, and FN is the number of false negatives. This metric varies from -1 to 1 , where 1 indicates an optimal solution, 0 a solution no better than a random guess, and -1 a total disagreement with the ground truth.

To enrich our evaluation of the network performance, we propose to use in addition three common evaluation metrics for binary classification problems, namely, *Accuracy*, *Precision*, and *Recall*, as defined in Equations (9)–(11).

$$Accuracy = \frac{TP + TN}{N} \quad (9)$$

where $N = TP + TN + FP + FN$ is the total number of samples.

$$Precision = \frac{TP}{TP + FP} \quad (10)$$

$$Recall = \frac{TP}{TP + FN} \quad (11)$$

If the *Accuracy* can give a false sense of performance of the model in a strongly imbalanced dataset (a naive model always classifying to the main class would have high *Accuracy*), its derivatives with respect to TP and TN are identical and constant. Changes in *Accuracy* are, therefore, easier to interpret than in other metrics.

4. Results

4.1. Results on the Test Dataset

We evaluated the predictions of the model on the test dataset, using the tool provided online by the hosting platform. Table 1 summarises the MCC and the rank reported, for the four proposed decision rules (cf. Section 3, (i)–(iv)), for $N_p = 200$, and for $N_p = 150$. Note that N_p represents the number of extracted pulses used as input.

Table 1. Results on the test dataset (evaluated online) for our four proposed decision rules and for N_p set to 200 and 150.

Settings	Kaggle VSB Test Data			
	$N_p = 150, w = 40$		$N_p = 200, w = 40$	
	MCC Private	Rank *	MCC Private	Rank *
(i) Baseline	0.580	1054	0.629	272
(ii) '1 = 3'	0.656	118	0.704	4
(iii) 1-Phase Opt.	0.631	245	0.658	65
(iv) 3-Phase (our)	0.665	46	0.704	4

* Ranking on the private leaderboard of Kaggle; (First place: MCC = 0.719).

The obtained results show that, first, our proposed approach provides state-of-the-art results. With settings (ii) and (iv), it achieves rank 4 on the leaderboard of the competition, and this is to our knowledge the best results achieved so far with a pure deep learning approach, without careful feature engineering. For comparison, the first place achieved an MCC of 0.719 (+2%), using a LightGBM model on carefully designed features using manually tuned threshold.

Second, the results demonstrate that considering $N_p = 200$ pulses per 50 Hz period leads to optimal results. This illustrates that the partial discharge pulses are not always the most dominant pulses in the signal. As the pulses are selected by amplitude, selecting more pulses increases the likelihood of capturing PD pulses. We also evaluated the performance for a larger number of pulses. However, the performance started to decrease, this may be explained either by the noise, or by the fact that increasing the input size also increases the number of parameters of the model and may lead to overfitting. The training dataset comprising around 8000 samples is rather small from a deep learning perspective.

4.2. Ablation Study

In addition to evaluating the performance of the proposed framework on the test dataset, we performed an ablation study to evaluate the impact of the different composing elements of the proposed network in more detail. As the test dataset does not contain any ground truth information, we cannot perform a more detailed analysis. Thus, we decided to split the training dataset into a new training dataset of size 6972 (6538 non-damaged and 434 damaged power-line samples), and a new test set of size 1740, (1599/141).

Table 2 summarizes the results of the ablation study. The evaluated experiments comprise the proposed model with the four decision rules (as discussed above) in part (1), and the ablated versions of this model. First, without the use of the Global Average Pooling layer ("w/o GAP") in part (2), second, without the last fully connected layer between the GAP layer and the output ("w/o FC") in part (3). We only present the results for $N_p = 200$ and $w = 40$ as this was the best performing parameter on the cross-validation.

Table 2. Ablation study and impact of the decision rules on the proposed model with $N_p = 200$ and $w = 40$.

	Methods	MCC	Acc.	Precision	Recall
(1)	(i) Baseline	0.710	0.959	0.772	0.695
	(ii) '1 = 3'	0.802	0.963	0.687	0.979
	(iii) 1-Phase Opt.	0.727	0.953	0.669	0.844
	(iv) 3-Phase (our)	0.817	0.967	0.726	0.957
(2)	w/o GAP-(i)	0.583	0.944	0.712	0.525
	w/o GAP-(ii)	0.722	0.952	0.656	0.851
	w/o GAP-(iii)	0.598	0.945	0.709	0.553
	w/o GAP-(iv)	0.616	0.955	0.683	0.610
(3)	w/o FC-(i)	0.643	0.952	0.784	0.567
	w/o FC-(ii)	0.751	0.959	0.702	0.851
	w/o FC-(iii)	0.701	0.954	0.700	0.745
	w/o FC-(iv)	0.775	0.966	0.772	0.816
	(STL + SVM) [†]	0.779 [▷]	0.963 [▷]	0.73 [▷] /0.68 [*]	0.88 [*]
	(LSTM) [‡]	0.344 [▷]	0.765 [▷]	0.23 [▷] /0.79 [*]	0.81 [*]

[†] Reported results on training set [12]. [‡] Reported results on samples drawn in the artificially augmented training set (split: non-PD: 807 / PD: 935) [18]. ^{*} Metrics as reported. [▷] Metrics recomputed on our own data split assuming constant sensitivity and specificity of the model.

5. Discussion

5.1. Results Analysis

The Global Threshold provides the best results: The results both of the ablation study and of the evaluation on the test dataset (Tables 1 and 2) show that setting (iv) always outperforms others both in terms of *MCC* and in terms of *Accuracy*. The only exception is the network without GAP layer. As this architecture performs worst overall, the comparison of the settings may not be very insightful in that case.

3-phase classification improves the performance: In general, the results demonstrate that considering the three phases together (settings (ii) versus (i) and (iv) versus (iii)) is always improving the overall performance (*MCC* and *Accuracy*), independently of the other parameters. This also makes sense from an application perspective as the essential information for the operator is to know whether a power line is damaged or not. Furthermore, some types of damage may impact all phases.

Optimising the threshold favors the *Recall* over the *Precision*: Looking at the *MCC* in all setups, it appears that decreasing the detection threshold to an optimized threshold improves the results (comparing the settings (iii) versus (i) and (iv) versus (ii)). This is in fact due to the nonlinearity of *MCC*, which favors the detection of the smallest class (the true positives, TP) strongly over the true negatives (TN). Comparing settings (iii) versus (i) in the part (1) of Table 2 illustrates the strong impact of the revised threshold on the *Recall*, while the *Precision* is harmed. From an application perspective, it seems indeed more important to detect true positives, even if it increases the number of false alarms. This is especially true in our case since a follow-up expert confirmation of positive cases is simplified thanks to the Pulse Activation Map presented in the next section. On the contrary, unnoticed damaged power-lines can lead to cascading damages and stronger consequences on the power distribution system.

Single-Phase Detection also provides a good performance: An additional interesting observation is that even though considering the detection results of three phases jointly provides the best performance, the proposed model is still able to provide good detection performance when the phases are considered independently (settings (i) and (iii) in Table 2). This is especially true for setting (iii). This confirms that the proposed framework can also be used on single phase measurements if all three phases are not available simultaneously.

The GAP layer is essential for the performance of the proposed architecture: The ablation study also confirms the significance of the GAP layer. Irrespective of the settings, the *MCC* always decreases when the GAP layer is removed from the architecture (part (2) of Table 2). The biggest impact of the GAP layer appears to be on the *Recall*, that is, the ability of the model to correctly identify true positives. As discussed in Section 2.2, GAP improves the generalization ability of the networks by decreasing the number of model parameters. For the network used here with $N_p = 200$ and $w = 40$, using the GAP layer instead of a flatten layer reduces the number of parameters from 518,113 to 6369. The resulting model is less prone to overfitting, which may explain its better performance. In addition, the GAP layer allows us to extract the Pulse Activation Map as described in Section 5.3.

The last FC layer improves the *Recall* significantly: Contrary to the architectures commonly encountered in the literature, adding a FC layer between the GAP and the output improves the results in our case study. This can be inferred by comparing part (1) and (3) of Table 2. The absence of the FC layer seems to harm the discriminative capacity of the network as the *Precision* increases slightly while the *Recall* decreases strongly. In fact, the model identifies more true negatives TN but far less TP. The *Accuracy* is in fact on par with the best models. However, the *MCC* strongly decreases. As already discussed above, from an application perspective, we believe that the *Recall* is more important than the *Precision*. This is in line with the competition objective of *MCC* maximization.

Comparison with the state-of-the-art: The proposed network outperforms state of the art methodologies found in the literature on both the *Recall* and the *MCC* and would be considered as superior according to the competition objective.

5.2. Note on the MCC for Imbalanced Dataset

The results in Table 2 indicate that higher *MCC* are always linked to higher *Recall* values, but not always to higher *Precision* values. This is in fact explained by the higher partial derivative of the *MCC* with respect to the true detection of the smallest class (here, true positives) compared to the detection of the largest class (here, true negatives). This is illustrated in Figure 2, where the *MCC* is represented using a colormap as the function of both TP and TN. On each side, a slice of the map *MCC* is represented when either the TP-rate or the TN-rate are, respectively, set to 95%. In Figure 2, the top sub-plot and rightmost sub-plot represent the corresponding partial derivatives. Excluding the region of low TN, the partial derivative of the *MCC* with respect to the TP is always much higher compared to the partial derivative with respect to the TN.

As only the *MCC* is reported for the test dataset by the online tool, it gives little insights on the impact of the different experimental setups on the TP and TN values. Changes in TP would dominate similar changes in TN. As discussed previously, the use of this metric is justified from an application perspective, as it favors *Recall* over *Precision*.

5.3. Pulse Activation Maps (PAM)

To provide more insights on the decision of the network and enable interpretability of the obtained results, we propose to use the PAM as per Equation (6) as a tool for a deeper analysis of the pulses for domain experts. Plots of some representative sample maps are demonstrated in Figure 3.

Potentially, the most valuable information for the domain experts is the evaluation of the activation of the pulses, as demonstrated in Figure 4, providing an information which parts of the pulses have contributed to the respective decision of the algorithm (either positive or negative).

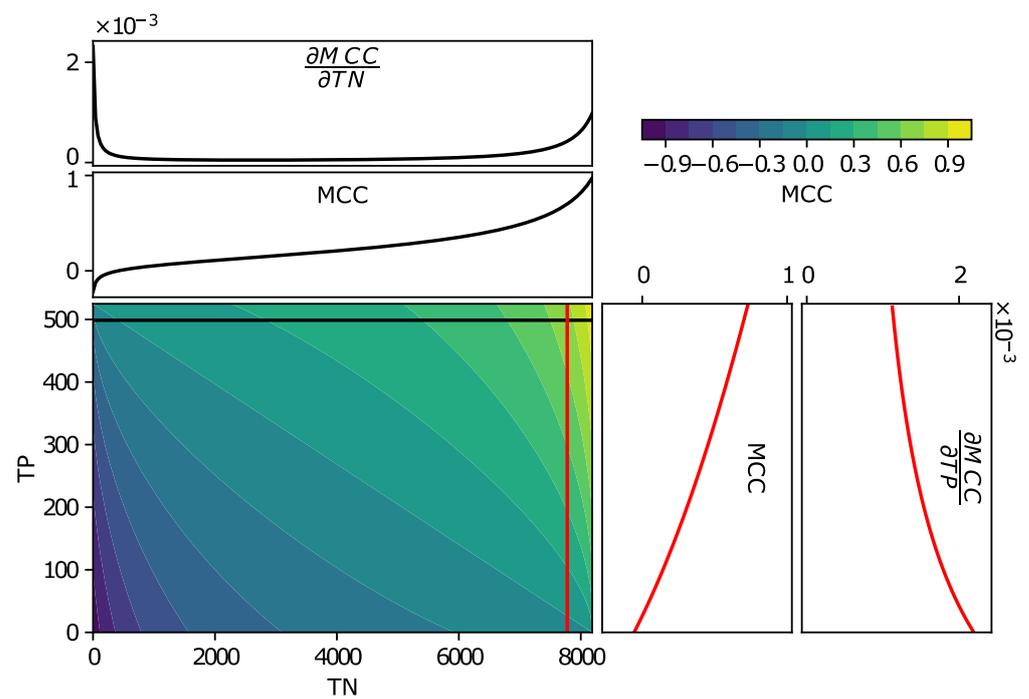


Figure 2. MCC: Colormap of the MCC. From colormap to top, $TP = 95\%$, MCC, and its derivative as a function of TN. From colormap to right, $TN = 95\%$, MCC, and its derivative as a function of TP.

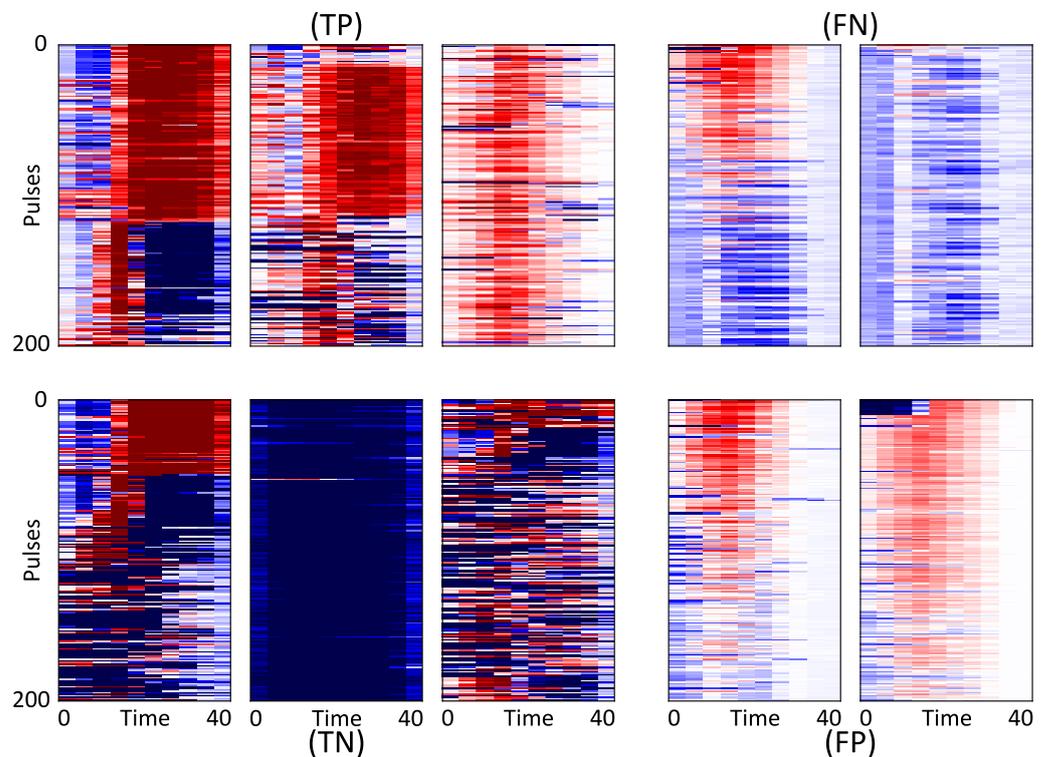


Figure 3. PAM: For TP (3 upper left), FN (2 upper right), TN (3 lower left), and FP (2 lower right). The color mapping is symmetric around 0 (white), from dark-blue (strongly negative) to dark-red (strongly positive). Note the the actual values matter little as the decision of the network only depends on whether the map is positive or negative in average.

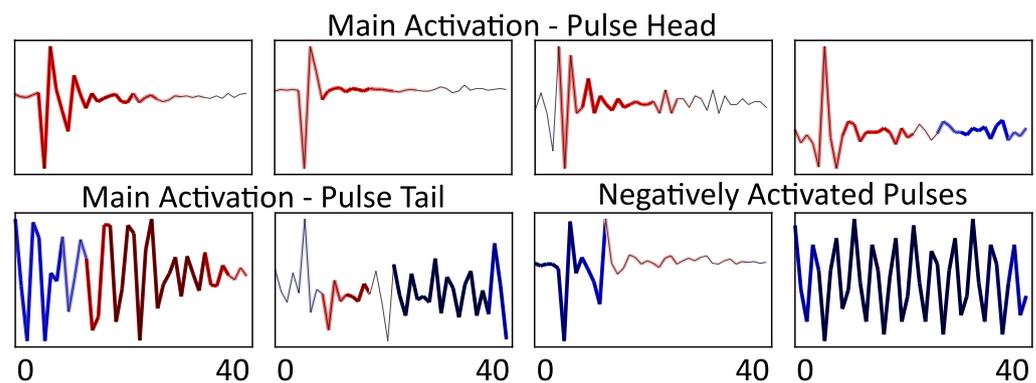


Figure 4. Pulse Activation: Some typical pulses with their activation (same color-map as in Figure 3). Upper row: pulses positively activated. Lower row left, the tail is rather considered as distinctive by the network. Lower row, right, strongly negatively activated pulses (the last one appears to be noise). The Y-axes are independent.

Interestingly, two types of patterns appear to influence the decision (cf. Figure 3): Samples for which the activation has higher amplitude for the pulse tails (e.g., first TP and first TN), and samples for which the pulse itself is the deciding factor (last TP, and the two displayed FP). These different types of pulses are presented in Figure 4. Remarkably, it appears that the decision to label one power line as damaged or not depends as much on the presence of positively activated parts as on the presence of pulses that are strongly deactivated (negative parts of the map). For example, the difference between the first TP and the first TN maps appears to rely on the number of pulses whose tails are positively activated. Furthermore, for the last FP, the largest pulses (the one at the top of the map) were identified as not being an indicator of a damaged power line. However, this is compensated on average by the rest of the map which is weakly positively activated. One could infer that the line may have some PD. However, those PD are not strong enough for the line to be considered as damaged. The first plotted FN in the figure can lead to a similar interpretation. The network did find some PD pulses in the signal (upper part is mostly red). However, they were apparently not sufficiently strong to compensate for the second part of the map which is weakly deactivated (light blue).

Finally, Figure 4 first row, presents pulses that are only positively activated. They could, therefore, be interpreted as representative of some typical PD pulses. An input from domain experts could provide additional insights and evaluate the match of the learnt patterns to the expert intuition on the type of the PD. This would guide additionally the interpretation of the obtained results.

6. Conclusions

In this paper, we proposed a new framework for the detection of damaged power lines. The proposed approach offers several improvements with respect to traditional power-line diagnostics. First, the proposed framework does not require any feature engineering and is able to handle raw measurements with extremely little preprocessing. Second, it provides competitive detection results at the power-line level, but also at the phase level. The proposed approach is robust and can detect damages in power lines from a single period of utility frequency. It provides a significant speed up compared to the more traditional PRPD approaches that require first, the processing of several hundreds of periods, and second, an expert analysis of the diagrams.

In addition, we proposed to extract the Pulse Activation Maps to improve the interpretability and to gain understanding on which part of the electrical signals are learned by the network as being a signature of a damaged power line. PAM can be used by the domain experts to gain more insights in the decisions of the proposed neural networks and to perform the diagnostics. PAM provides the information on which pulses and which

part of the pulses dominated the decision of the neural network and allows to verify the network's decision.

It can be pointed out that one limit of the task we tackled here is the relatively small size of the training dataset (from a deep learning perspective). Even though very competitive results were obtained, we believe our approach can showcase its full potential when more and more data will be available. Training the framework with more data will allow for a more precise tuning of the hyperparameters. Furthermore, if samples were identified per power-line, timestamped and collected over a long time period (which was not the case in the considered case study), the monitoring of the PAM evolution over time would be a very promising follow-up research. We could expect that, as a power-line damage increases, the PAM would become more and more positively activated, and such monitoring would have a potentially large benefit for the utility operators.

Author Contributions: Conceptualization, G.M. and O.F.; Formal analysis, G.M. and C.-C.H.; Funding acquisition, O.F.; Investigation, G.M. and C.-C.H.; Methodology, G.M., C.-C.H. and O.F.; Project administration, G.M. and O.F.; Resources, O.F.; Software, G.M. and C.-C.H.; Supervision, G.M. and O.F.; Validation, G.M. and O.F.; Visualization, G.M. and C.-C.H.; Writing—original draft, G.M. and C.-C.H.; Writing—review & editing, G.M. and O.F.; All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Innosuisse grant no. 27662.1 PFES-ES.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The VSB Power Line Fault Detection dataset of the ENET Centre is openly available <https://www.kaggle.com/c/vsb-power-line-fault-detection> (accessed on 4 July 2020).

Conflicts of Interest: The authors declare no conflict of interest.

References

- Chen, Y.Q.; Fink, O.; Sansavini, G. Combined fault location and classification for power transmission lines fault diagnosis with integrated feature extraction. *IEEE Trans. Ind. Electron.* **2017**, *65*, 561–569. [CrossRef]
- Mashikian, M.; Szarkowski, A. Medium voltage cable defects revealed by off-line partial discharge testing at power frequency. *IEEE Electr. Insul. Mag.* **2006**, *22*, 24–32. [CrossRef]
- Kawaguchi, Y.; Yanabu, S. Partial-discharge measurement on high-voltage power transformers. *IEEE Trans. Power Appar. Syst.* **1969**, *88*, 1187–1194. [CrossRef]
- Ibrahim, K.; Sharkawy, R.; Salama, M.; Bartnikas, R. Realization of partial discharge signals in transformer oils utilizing advanced computational techniques. *IEEE Trans. Dielectr. Electr. Insul.* **2012**, *19*, 1971–1981. [CrossRef]
- Khan, Q.; Refaat, S.S.; Abu-Rub, H.; Toliyat, H.A. Partial discharge detection and diagnosis in gas insulated switchgear: State of the art. *IEEE Electr. Insul. Mag.* **2019**, *35*, 16–33. [CrossRef]
- McGreevy, D.; Giussani, R.; Seltzer-Grant, M.; Singh, A.; Patel, A.; Calladine, S.; Gibb, G. Deployment of an online partial discharge monitoring system for power station with focus on gas turbine generators. In Proceedings of the 2017 INSUCON-13th International Electrical Insulation Conference, Birmingham, UK, 16–18 May 2017; pp. 1–5.
- Misák, S.; Fulnecek, J.; Vantuch, T.; Buriánek, T.; Jezowicz, T. A complex classification approach of partial discharges from covered conductors in real environment. *IEEE Trans. Dielectr. Electr. Insul.* **2017**, *24*, 1097–1104. [CrossRef]
- Hudon, C.; Belec, M. Partial discharge signal interpretation for generator diagnostics. *IEEE Trans. Dielectr. Electr. Insul.* **2005**, *12*, 297–319. [CrossRef]
- Strachan, S.; Rudd, S.; McArthur, S.; Judd, M.; Meijer, S.; Gulski, E. Knowledge-based diagnosis of partial discharges in power transformers. *IEEE Trans. Dielectr. Electr. Insul.* **2008**, *15*, 259–268. [CrossRef]
- Karami, H.; Tabarsa, H.; Gharehpetian, G.B.; Norouzi, Y.; Hejazi, M.A. Feasibility Study on Simultaneous Detection of Partial Discharge and Axial Displacement of HV Transformer Winding Using Electromagnetic Waves. *IEEE Trans. Ind. Inform.* **2019**, *16*, 67–76. [CrossRef]
- Raymond, W.J.K.; Ilias, H.A.; Abu Bakar, A.H. Classification of partial discharge measured under different levels of noise contamination. *PLoS ONE* **2017**, *12*, e0170111.
- Dong, M.; Sun, Z.; Wang, C. A pattern recognition method for partial discharge detection on insulated overhead conductors. In Proceedings of the 2019 IEEE Canadian Conference of Electrical and Computer Engineering (CCECE), Edmonton, AB, Canada, 5–8 May 2019; pp. 1–4.

13. Karimi, M.; Majidi, M.; MirSaeedi, H.; Arefi, M.M.; Oskuoee, M. A Novel Application of Deep Belief Networks in Learning Partial Discharge Patterns for Classifying Corona, Surface, and Internal Discharges. *IEEE Trans. Ind. Electron.* **2019**, *67*, 3277–3287. [[CrossRef](#)]
14. Li, G.; Rong, M.; Wang, X.; Li, X.; Li, Y. Partial discharge patterns recognition with deep Convolutional Neural Networks. In Proceedings of the 2016 International Conference on Condition Monitoring and Diagnosis (CMD), Xi'an, China, 25–28 September 2016; pp. 324–327.
15. Banno, K.; Nakamura, Y.; Fujii, Y.; Takano, T. Partial Discharge Source Classification for Switchgears with Transient Earth Voltage Sensor Using Convolutional Neural Network. In Proceedings of the 2018 Condition Monitoring and Diagnosis (CMD), Perth, Australia, 23–26 September 2018; pp. 1–5.
16. Wang, G.; Yang, F.; Peng, X.; Wu, Y.; Liu, T.; Li, Z. Partial discharge pattern recognition of high voltage cables based on the stacked denoising autoencoder method. In Proceedings of the 2018 International Conference on Power System Technology (POWERCON), Guangzhou, China, 6–9 November 2018; pp. 3778–3792.
17. Nguyen, M.T.; Nguyen, V.H.; Yun, S.J.; Kim, Y.H. Recurrent neural network for partial discharge diagnosis in gas-insulated switchgear. *Energies* **2018**, *11*, 1202. [[CrossRef](#)]
18. Dong, M.; Sun, J. Partial discharge detection on aerial covered conductors using time-series decomposition and long short-term memory network. *Electr. Power Syst. Res.* **2020**, *184*, 106318. [[CrossRef](#)]
19. Qu, N.; Li, Z.; Zuo, J.; Chen, J. Fault Detection on Insulated Overhead Conductors Based on DWT-LSTM and Partial Discharge. *IEEE Access* **2020**, *8*, 87060–87070. [[CrossRef](#)]
20. Wu, M.; Cao, H.; Cao, J.; Nguyen, H.L.; Gomes, J.B.; Krishnaswamy, S.P. An overview of state-of-the-art partial discharge analysis techniques for condition monitoring. *IEEE Electr. Insul. Mag.* **2015**, *31*, 22–35. [[CrossRef](#)]
21. Barrios, S.; Buldain, D.; Comech, M.P.; Gilbert, I.; Orue, I. Partial discharge classification using deep learning methods—Survey of recent progress. *Energies* **2019**, *12*, 2485. [[CrossRef](#)]
22. Acharya, U.R.; Oh, S.L.; Hagiwara, Y.; Tan, J.H.; Adam, M.; Gertych, A.; San Tan, R. A deep convolutional neural network model to classify heartbeats. *Comput. Biol. Med.* **2017**, *89*, 389–396. [[CrossRef](#)]
23. Susto, G.A.; Cenedese, A.; Terzi, M. Time-Series Classification Methods: Review and Applications to Power Systems Data. In *Big Data Application in Power Systems*; Elsevier: Amsterdam, The Netherlands, 2018; pp. 179–220.
24. Fawaz, H.I.; Forestier, G.; Weber, J.; Idoumghar, L.; Muller, P.A. Deep learning for time series classification: A review. *Data Min. Knowl. Discov.* **2019**, *33*, 917–963. [[CrossRef](#)]
25. Zeng, Q.; Pan, H.; Chen, B.; Liao, Z. Research on STLF Method Based on One-Dimensional Convolution and Slope Feature. In Proceedings of the 2019 IEEE International Conference on Power Data Science (ICPDS), Taizhou, China, 12–15 December 2019; pp. 118–121.
26. Habeeb, R.A.A.; Nasaruddin, F.; Gani, A.; Hashem, I.A.T.; Ahmed, E.; Imran, M. Real-Time Big Data Processing for Anomaly Detection: A Survey. *Int. J. Inf. Manag.* **2019**, *45*, 289–307. [[CrossRef](#)]
27. Kuo, C.C.J. Understanding convolutional neural networks with a mathematical model. *J. Vis. Commun. Image Represent.* **2016**, *41*, 406–413. [[CrossRef](#)]
28. ENET Centre at VSB—Technical University of Ostrava. VSB Power Line Fault Detection. 2019. Available online: <https://www.kaggle.com/c/vsb-power-line-fault-detection> (accessed on 4 July 2020).
29. Butterworth, S. On the theory of filter amplifiers. *Wirel. Eng.* **1930**, *7*, 536–541.
30. Simonyan, K.; Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv* **2014**, arXiv:1409.1556.
31. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
32. Lin, M.; Chen, Q.; Yan, S. Network in network. *arXiv* **2013**, arXiv:1312.4400.
33. Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; Torralba, A. Learning deep features for discriminative localization. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 2921–2929.
34. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
35. Matthews, B. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochim. Biophys. Acta BBA Protein Struct.* **1975**, *405*, 442–451. [[CrossRef](#)]