This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Tervo, Roope; Karjalainen, Joonas; Jung, Alexander
Predicting Electricity Outages Caused by Convective Storms

Published in:
2018 IEEE Data Science Workshop, DSW 2018 - Proceedings

DOI:
[10.1109/DSW.2018.8439906](https://doi.org/10.1109/DSW.2018.8439906)

Published: 17/08/2018

Document Version
Peer reviewed version

Please cite the original version:
Tervo, R., Karjalainen, J., & Jung, A. (2018). Predicting Electricity Outages Caused by Convective Storms. In *2018 IEEE Data Science Workshop, DSW 2018 - Proceedings* (pp. 145-149). [8439906] IEEE.
<https://doi.org/10.1109/DSW.2018.8439906>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

©Copyright 2018 IEEE. Published in the 2018 IEEE Data Science Workshop (DSW 2018), scheduled for June 4-6, 2018 in Lausanne, Switzerland. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the IEEE. Contact: Manager, Copyrights and Permissions / IEEE Service Center / 445 Hoes Lane / P.O. Box 1331 / Piscataway, NJ 08855-1331, USA. Telephone: + Intl. 908-562-3966.

PREDICTING ELECTRICITY OUTAGES CAUSED BY CONVECTIVE STORMS

Roope Tervo, Joonas Karjalainen

Observing and information systems centre
Finnish Meteorological Institute
B.O. 503, 00101 Helsinki, Finland

Alexander Jung

Dept of Computer Science
Aalto University
B.O. 11000, 00076 Aalto, Finland

ABSTRACT

We consider the problem of predicting power outages in an electrical power grid due to hazards produced by convective storms. These storms produce extreme weather phenomena such as intense wind, tornadoes and lightning over a small area. In this paper, we discuss the application of state-of-the-art machine learning techniques, such as random forest classifiers and deep neural networks, to predict the amount of damage caused by storms. We cast this application as a classification problem where the goal is to classify storm cells into a finite number of classes, each corresponding to a certain amount of expected damage. The classification method uses as input features estimates for storm cell location and movement which has to be extracted from the raw data.

A main challenge of this application is that the training data is heavily imbalanced as the occurrence of extreme weather events is rare. In order to address this issue, we applied SMOTE technique.

Index Terms— Power distribution, Weather Impact, Random Forest Classifier, Multilayer Perceptron, Neural Network

1. INTRODUCTION

A key problem faced by operators of power grids is the prediction of damages caused by extreme weather events. The hazards produced by intense winds, lightning and tornadoes have significant social impacts and cause remarkable liability for damages for distribution companies. The main objective of this work is to apply machine learning techniques in order to obtain a real-time prediction of the short-term damage potential based on current weather information.

As the weather-caused damages incur significant economic loss, a lot of effort has been put into studying efficient prediction methods. The problem of storm cell identification and tracking has been studied thoroughly in [1]. We will use the methods developed in [1] as pre-processing of raw weather data. The authors of [2] provide an excellent overview of several successful projects. In general, random forest classifiers (RFC) [3] are very popular method for analysing weather impacts. The authors of [4] applied a

variant of AdaBoost⁺ to predict outages caused by wind and lightning in overhead distribution systems.

This paper is organised as follows: In Section 2, we formalise the problem of predicting outages as a classification problem using estimates of storm cell location and movements as input features. We then discuss the application of RFC and deep neural network classifiers in Section 3. Some illustrative numerical experiments based on historic data collected by the Finnish Meteorological Institute (FMI) are discussed in Section 4.

2. PROBLEM SETUP

We identify convective storm cells by contouring weather radar reflectivity composite CAPPI (constant altitude plan position indicator) images with a solid 35 DBZ threshold. After contouring, we apply the DBSCAN method [5] to cluster the contoured objects to an identifiable storm object which are stored to a PostGIS database.

Within the DBSCAN method, a storm object is considered as a core point if the area sum of its nearby storm objects exceeds the given area threshold. The storm objects that are within the neighbourhood of a core point but do not fulfil the criteria mentioned above, are considered as outliers. Together with their outliers, connected core points form a cluster and storm objects that belong to neither of these classes are regarded as noise. In this study, the area limit was set to 20km² and the neighbourhood radius was set to 2km². After clustering, an object-oriented storm tracking algorithm using optical flow [6] is used to track and forecast paths of the storms. The storm cell tracking is done for a time horizon of 2 hours ahead with a time resolution of 5 minutes.

In order to recognise damage potential of convective storms and to predict future outages, we categorise storm cells into four classes based on how much damage they are expected to cause for the power grid. In more detail, the storm cells are assigned to a class based on how large share of transformers under the storm are without electricity. The damage may happen to any point in the grid, but outages are always reported at the transformer nodes of the power grid.

Thus, we formulate the outage prediction as a classifica-

tion problem which aims at classifying storm cells into one of four different classes, described in table 1. The particular choice of this classes aims to provide a simple 'at glance' view which is convenient for the end user (power grid operator).

Table 1: Class definitions of the storm cells.

Class	Share of transformers
0	no damage
1	0 - 10 %
2	10 - 50 %
3	50 - 100 %

We used data collected by FMI during years 2012 to 2017 to train the classification methods. This data contains identified storm cell objects at a time resolution of five minutes. Note that while most of the time steps do not contain any storm cells, during "active" days, each time step may contain hundreds of cells. Fetching all storm cell objects covering the power grid at some point of their life cycle provided 886 020 training samples.

Each training sample represents one individual storm cell which is characterized by the list of features listed in Table 2. The data is very imbalanced as most of the convective cells are not powerful enough to cause harm for the power grid. We depict a histogram of the target classes contained in the training dataset in Figure 3. In particular, we have 551 029 samples of class 0 (no harm), 4919 samples of class 1, 4286 samples belonged to class 2 and only 3337 belong to class 3 (most harmful).

Many ground observations were missing as they were fetched from the nearest observation station to the storm center and capabilities of those weather stations varies a lot. Absent parameters were initialised to zero to ensure coherence of the data. Effects of this inadequate data is discussed more in section 4.

Outage data and power grid description are fetched from two power distribution companies. Spatial coverage of the data is shown in image 1. The data set contains in total 33 858 outages.

3. THE CLASSIFICATION METHOD

We created two alternative methods for classification. First, we used RFC to get 'baseline' performance. The classifier was implemented without limitations in tree size and with equal class weights. For the training of the RFC, we used the Gini impurity as loss function, i.e.

$$G = - \sum_{i=1}^{n_c} (p_i(1 - p_i)) \quad (1)$$

where n_c is the number of classes and p_i is the share of i^{th} class in the tree.

Table 2: Used input features

Feature	Explanation
Area	Area covered by the storm cell
Age	Age of the storm
Lightning density	Lightning density under storm cell
Max DBZ	Maximum radar reflectivity of the storm cell (spatially). Represents maximum rain intensity.
Min DBZ	Minimum radar reflectivity of the storm cell (spatially). Represents minimum rain intensity.
Mean DBZ	Mean radar reflectivity of the storm cell (spatially)
Median DBZ	Median radar reflectivity of the storm cell (spatially)
Std of DBZ	Standard deviation of radar reflectivity of the storm cell (spatially)
Lat	Storm center latitude
Lon	Storm center longitude
Temperature observations	Air temperature from ground
Pressure observations	Air pressure from ground
Wind speed observations	Wind speed from ground
Wind direction observations	Wind direction from ground
Precipitation amount	Precipitation amount from ground observations
Snow depth observations	Snow depth from ground

As an alternative to RFC, we also implemented a classifier based on a multi layer perceptron (MLP) neural network [7]. The network structure is described in Figure 2. The first layer contains 20 nodes wide dense layer. In the following layers, the number of nodes is reduced to 16, 8, 4 and finally to 1 node. The first three dense layers use the rectified linear unit (Relu) as activation and dropout regularisation layers are included after first and second dense layers. In the final layer, we used the "Softmax" activation function in order to obtain the predicted class probabilities in the output layer.

In order to cope with the imbalanced training data we decided to use the cross entropy loss function [8]. This loss is defined as

$$H(p, q) = - \sum_x (p(x) \log(q(x))) \quad (2)$$

where $p(x)$ is a probability distribution of true labels and $q(x)$ is a probability distribution of predicted labels. Categorical entropy is a good default choice and it has an advantage that different classes can be easily preferred by giving different weights for the classes.

Fig. 1: Spatial coverage of power grid information available in this project. Soft lines represents municipality districts and darker lines power grids. Image is from Southern Finland.

Fig. 2: Network structure for classification task

We used the Adam optimiser [9] for training the model to avoid challenging points in the optimisation space. All hyper parameters are shown in the table 3.

4. EXPERIMENTS

We divided the data set into training and validation set with share of 75 % and 25 % respectively. The RFC worked relatively well without requiring any tuning of hyper parameters. The classification accuracy obtained for the training set was around 98 % and for the validation set up to 88 %. Thus, the RFC tends to slightly over-fit the training data.

The MLP network required significantly more efforts for tuning the hyperparameters (see Table 3). Moreover, and in order to balance the training data, we combined the MLP classifiers with the synthetic minority over-sampling technique

Table 3: Hyperparameters of the MLP classifier

Parameter	Value
Batch size	256
Epoch count	1000
Dropout probability	10 %
α (learning rate)	0.001
β_1 (exp decay for momentum)	0.9
β_2 (exp decay for momentum)	0.999
ϵ (stability constant)	10^{-8}
Initial decay	no decay

Fig. 3: Histogram of classes in original dataset, after generating new samples synthetically and in validation dataset. The original dataset is very imbalanced but after generating synthetic samples by SMOTE, training dataset is completely balanced.

(SMOTE) [10]. The method generates new training samples in the vicinity of the original training samples by interpolating their $k = 5$ nearest neighbours (in the feature space) as following:

$$x_{new} = x_i + \lambda \times (x_{zi} - x_i) \quad (3)$$

where x_i is an original minority class sample, x_{zi} is one of x_i 's k nearest neighbour and λ is random variable drawn uniformly from the interval $[0, 1]$. The synthetic data set generated by SMOTE contains data points with a balanced distribution of classes (see figure 3).

Batch size was compromised to provide enough performance in sufficient time. Amount of hidden layers, size of hidden layers and learning rate was searched by trial and error. Training and validation loss and accuracy are plotted in figure 4. There is no sign of over-fitting and thus quite low dropout probability was used for both dropout layers [8].

As mentioned in section 2, the data contained large amount of incomplete samples. Intuitive assumption would say, that filtering those samples would be beneficial to gain better results. We created a new dataset X_{fit} from samples which contained all parameters and used that to train the classification methods. The new dataset contained 563 571 samples. Filtering clean samples from the dataset did not improve the results.

Fig. 4: Training and validation metrics while training the MLP network. Validation performance is significantly better than training performance since training dataset contains synthetically generated samples which do not exist in the validation set.

5. RESULTS

The results are shown in table 4. In the end, RFC performed better for predicting amount of damage. It is notable that AUC, loss or accuracy do not catch differences very well as class 0 dominates all metrics. Thus differences between methods can be best shown in confusion matrices (figure 5 and 6).

A performance of RFC is excellent for class 0 (no damage) with over 99 % accuracy and good for class 3 (most harmful) with 87 % accuracy. The classifier has little problems with distinguishing classes 1 and 2 as it tends to underestimate a damage potential of storms. The results of MLP classifier was significantly worse. While MLP provided 90 % accuracy for class 0 and 75 % accuracy for class 3, accuracy for classes 1 and 2 are only near 50 %.

Classification is fast with both classifiers. For example, 221 506 samples can be classified in about 5 second with two 3,3 GHz i5 CPUs. While optical flow is a lightweight algorithm, DBSCAN method has a $\mathcal{O}(n \log n)$ computational cost [5], which may be challenging during “active days”.

Table 4: Metrics for different methods. ‘MLP SMOTE’ means MLP neural network with synthetic samples and ‘MLP X_{flt} ’ means MLP with synthetic samples without samples with missing values. Micro average for F1 score is calculated by counting the total true positives, false negatives and false positives over all classes.

Metrics	MLP SMOTE	MLP X_{flt}	RFC
AUC	0.96	0.95	0.99
Validation accuracy	89 %	85 %	96 %
F1 score			
micro average	71 %	85 %	99 %

Fig. 5: Confusion matrix of Random Forest Classifier

Fig. 6: Confusion matrix of MLP

6. CONCLUSIONS

This paper studied the application of RFC and MLP classifiers to the problem of predicting power grid outages caused by hazardous storm cells. Some illustrative numerical experiments based on weather data collected by FMI indicated that RFC can outperform deep MLP in predicting the amount of damage caused by storm cells.

This work suggests several interesting avenues for future research. One promising direction is to use more advanced models and methods for the training data, e.g., times series models and recurrent neural networks. So far, we also used only very basic methods for coping with missing data (just replace by zero) and imbalanced training data (using SMOTE). It would be interesting to apply more advanced techniques for coping with imbalanced data, e.g., the “Rare-Transfer” algorithm [11].

Currently Random Forest Classifier is used in operational application.

7. REFERENCES

- [1] Pekka Juhana Rossi, *Object-Oriented Analysis and Nowcasting of Convective Storms in Finland*, Ph.D. thesis, Aalto University, 2015.
- [2] Amy McGovern, Kimberly L. Elmore, David John Gagne II, Sue Ellen Haupt, Christopher D. Karstens, Ryan Lagerquist, Travis Smith, and John K. Williams, “Using artificial intelligence to improve real-time decision-making for high-impact weather,” *Bulletin of the American Meteorological Society*, vol. 98, no. 10, pp. 2073–2090, 2017.
- [3] Leo Breiman, “Random forests,” *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [4] Padmavathy Kankanala, Sanjoy Das, and Anil Pahwa, “Adaboost⁺: An ensemble learning approach for estimating weather-related outages in distribution systems,” *IEEE Transactions on Power Systems*, vol. 29, no. 1, pp. 359–367, 2014.
- [5] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al., “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *KDD-96 Proceedings*, 1996, vol. 96, pp. 226–231.
- [6] Berthold KP Horn and Brian G Schunck, “Determining optical flow,” *Artificial intelligence*, vol. 17, no. 1-3, pp. 185–203, 1981.
- [7] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, pp. 164–223, MIT Press, 2016, <http://www.deeplearningbook.org>.
- [8] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, pp. 255–265, MIT Press, 2016, <http://www.deeplearningbook.org>.
- [9] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.
- [10] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer, “Smote: synthetic minority over-sampling technique,” *Journal of artificial intelligence research*, vol. 16, pp. 321–357, 2002.
- [11] Samir Al-Stouhi and Chandan K Reddy, “Transfer learning for class imbalance problems with inadequate data,” *Knowledge and information systems*, vol. 48, no. 1, pp. 201–228, 2016.