

Strategies for Intelligent Low-Voltage Network Monitoring – Detection of Unregistered Electric Vehicles Using a Recurrent Neural Network

Benjamin James Groene, Sabrina Hempel, Eckehard Tröster
Energynautics GmbH
 Darmstadt, Germany
 Email: bjg2463@gmail.com

Abstract—Due to the expected rise in electric vehicle (EV) penetration, distribution system operators are looking for ways to monitor the impact on the distribution system. The addition of load from EVs on a distribution system could lead to reliability issues, which may go unnoticed until they cause a major problem, as well as to unexpected voltage and thermal violations due to an increased peak load. The goal of monitoring is to shed light on any existing or future concerns. Without a monitoring solution, if a problem does become apparent, the cause would most likely go undiagnosed. Therefore, a machine learning solution to EV detection has been offered which detects EV charging events on a given feeder which occurred during the previous day. DSOs can then use this information to get a better idea of how many EVs are connected to a feeder as well as when they are typically being connected to the utility. EV information can then be used in power system calculations to yield more realistic results for the basis of project creation, greatly reducing the risk associated with EVs on the distribution system.

Keywords— distribution grid monitoring, impact of electric vehicles on grid, detection of electric vehicles using machine learning, distribution grid reliability, low-voltage network reliability, low-voltage grid reliability

I. INTRODUCTION

The notion of powering motor vehicles using electricity is becoming increasingly accepted as overall awareness and concern for the environment continues to grow in consumers. Thanks to the steadily rising share of renewable energy sources within the overall energy mix, powering vehicles via the power grid is becoming a more and more environmentally friendly option. Technological advances in the electromobility sector, expanding charging infrastructure, monetary incentives and decreasing battery costs are also among the key factors in the rise of electric vehicle (EV) popularity as previously hesitant consumers gradually begin seeing EVs as a viable means of daily transportation. Germany and many other countries are setting ambitious goals and investing large amounts of money into research and development, the expansion of charging infrastructure and various incentive programs with the aim of reaching higher EV penetrations. The goal in Germany is to reach one million EVs on the road by 2020. Experts predict this number to be reached but not until mid-2021 or 2022 [1]. The forecasted projection of accumulated new EV registrations through 2025 in Germany can be seen in Fig. 1. Distribution systems today were not designed with the intention of handling a large penetration of EVs and could therefore potentially face issues while doing so, especially in remote areas and on lines which were already heavily loaded prior to the addition of EV load.

Due to a lack of monitoring devices on the low-voltage grid coupled with uncertainty of how many EVs are being charged on a given feeder, distribution system operators (DSO) often fear that problems could already exist on the distribution grid or that problems could arise in the near future. Potential concerns could be, but are not limited to, unacceptable voltage levels at distribution system nodes and consequently across utility and customer equipment, thermal violations on distribution lines along with unacceptable equipment loading at utility and customer facilities, loss of reliability due to constricted load transferability during contingency events, load unbalance if EV load is not evenly distributed between the three phases and increased harmonic levels. Since the voltage and thermal issues are a product of an increased peak load, these issues are significantly more problematic in the case of uncontrolled charging, i.e. without utilization of demand-side management. For the reasons above, along with increasing penetration of distributed generation (DG), there has become a growing need for greater visibility in the low-voltage grid. It has thus become essential to monitor EV charging and its impacts on the distribution system and to use this information in power system calculations in order to obtain more realistic results for the basis of project creation. This is essential to reduce the risk associated with EVs on the distribution system. An electric vehicle detection solution will therefore be offered using deep learning in order to learn the behavior of a distribution feeder and detect when an electric vehicle has been plugged in. The impact of EVs has become a major topic of interest for power system engineers and DSOs. This paper is meant to deliver an initial machine learning solution to EV charging detection which may be further developed along with other projects aiming for total visibility of the distribution system.

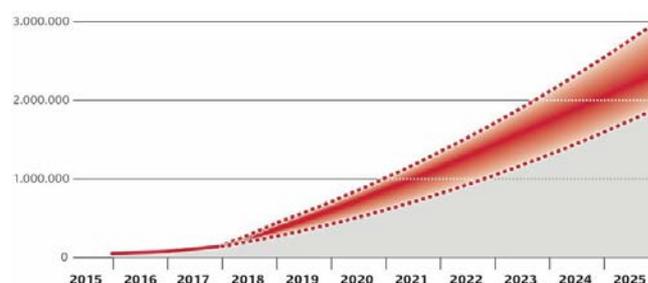


Fig. 1. Projection of Accumulated New EV Registrations through 2025 in Germany [1].

II. SMART METER ROLLOUT IN GERMANY

The smart meter rollout in Germany is underway and will open many new possibilities for the DSO to monitor the

distribution grid. The old Ferraris meters will be replaced with modern meters according to §29 Messstellenbetriebsgesetz (MsbG), which outlines which customers will be required to have a modern metering system (mME) installed and which customers will be required to have an intelligent metering system (iMSys) installed. An mME will not be equipped with a smart meter gateway (SMGW) and will therefore have no ability to transmit real-time data. Meter readings will be saved and stored on-site for a rolling two-year period. An iMSys is an mME or multiple mMEs connected through a SMGW and transmits meter readings to the Measurement Site Operator (Messstellenbetreiber) in fifteen-minute resolution, where it will then be processed and securely transmitted to the appropriate parties [2]. An illustration showing the customers who will be required to install an iMSys as well as the planned schedule for the smart meter rollout is provided in [3]. Consumers with EVs will only be required to install an iMSys if they are in accordance with §14a Energiewirtschaftsgesetz (EnGW). This means they have signed a contract with the DSO to be considered a controllable load in order to receive a reduced network charge. DSOs will only have access to monthly consumption and peak load data for each month according to §66 paragraph 1 MsbG and only have access to fifteen-minute meter readings under special circumstances outlined in §66 paragraph 2 MsbG. It is also possible for an EV to push a consumer over 6,000 kWh/a, which would require the consumer to replace their current meter with an iMSys. According to [3], the EV would not be metered separately and the DSO would only receive yearly consumption data. Additionally, §31 paragraph 4 MsbG states that yearly electricity consumption will be calculated using the average of the past three annual consumption values at a given metering point. This means that even if a consumer purchases an EV which causes their consumption to exceed 6,000 kWh/a, it could take up to three years before they would be required to install an iMSys. This lag time represents a significant gap in visibility which is another key driver for the solution offered in this paper. Therefore, it is improbable that the rollout of SMGWs will provide DSOs with enough data to accurately monitor EV charging alone, although it will improve the overall visibility of the distribution system and open up new possibilities for monitoring solutions in the future.

III. EXPERIMENTAL METHOD

An AI-based solution to EV detection has been developed using a supervised learning algorithm. The algorithm trains a bidirectional recurrent neural network (BRNN) equipped with long short-term memory (LSTM) cells what EV charging will look like so that when introduced to new data, the algorithm is able to differentiate between EV charging and other non-EV-charging events. The only data provided to the algorithm are current measurements taken at the secondary substation. For a deeper understanding of what supervised and unsupervised learning are, refer to [4]. To gain a first intuition of how recurrent neural networks (RNN) and LSTM cells function, refer to [5], and for a more technical understanding of RNNs and LSTM cells refer to [6]. RNNs are the state-of-the-art method in machine learning when dealing with time-series data and LSTM cells make it possible for the network to remember long-term

dependencies within lengthy sequences without experiencing the vanishing gradient problem which basic neural cells are highly prone to. For a technical understanding of vanishing and exploding gradients, refer to [7]. The additional bidirectional characteristic allows the algorithm to predict EV charging at a given timestep using dependencies from both past and future timesteps whereas standard unidirectional RNNs can only use information from timesteps occurring before the timestep in question. This component is essential for this task as a unidirectional RNN will not be sufficient to achieve high performance in this case. The data available for this experiment includes monthly current data in 500-millisecond resolution from two different feeders in Germany. It is unknown if there were already any EVs which were connected to these feeders, but for the purpose of this study it is assumed that there were not. The meaning of true positives, false positives and false negatives in this use case is illustrated in Fig. 2. Where, “Positive” corresponds to an EV charging event and “Negative” corresponds to a non-EV-charging event.

		Actual	
		Positive	Negative
Predicted	Positive	True Positive	False Positive
	Negative	False Negative	True Negative

Fig. 2. Confusion matrix for possible prediction outcomes [8].

Due to the lack of actual measurement data from a feeder with known EVs and known charging times, the available data was used as background data into which artificial EV charging events were injected. The approach was to inject charging events into all the data and then split the data into training, cross-development, and test sets. This tells the algorithm explicitly what EV charging will look like and when it occurs during training. The trained model is then used to detect charging in previously unseen data, i.e. the test set. Refer to [9] for a deeper understanding of training, cross-validation, and test sets and how they should be used and structured. This method requires the EV charging profiles given to the model during training to match the EV charging profiles which the model will see during testing and after project deployment. Therefore, it is important that the artificial charging events look very similar to how real charging events will appear on a feeder. The EV detection model will look at each phase individually so that each phase can be monitored separately. It can then be easily concluded whether the charging was single-phase, two-phase, or three-phase by simply checking how many phases EV charging was detected on at the timestep in question. EVs with charging amplitudes less than 15 A were not initially considered in the detection algorithm. It was assumed that the EVs with higher charging amplitudes would cause a larger threat to system reliability and could be detected with a much higher degree of certainty. Below 15 A, it was predicted that too many other loads operating in the same range would hamper performance. However, using the final model to detect EV charging events lower than 15 A was explored later, but only after the algorithm was proven effective for predicting charging events equal to or greater than 15 A. Two artificial

EV charging events randomly chosen in the range of 15 to 33 A were inserted at random times each day. This range was chosen considering common maximum EV charging power values. This range includes single-phase charging from 3.6 to 7.4 kW and three-phase charging from 11 to 22 kW. This method provides varying EV charging amplitudes and forces the algorithm to be open for EV detection at any hour of the day. It will also provide an analysis tool to compare the performance of the model at different hours of the day. EVs require high current levels to charge, but what really sets them apart from other loads is the duration in which they require these high current levels, i.e. normally for at least one hour. Looking at some charging data from [10], it can be seen that full current consumption is almost always reached within 10 seconds. This is not the case with the end of the charging cycle, which is slowly tapered until the battery is finished charging (unless the EV is unplugged before the end of the charging cycle). For this reason, the 500-millisecond data was transformed into 10-second resolution to eliminate the transient portion and the ramping portion of the charging profile. This also acts to shorten the sequences by a factor of 20, which decreases training time and memory consumption during training drastically. The 10-second data was then converted to delta values to force charging events to appear as a current spike. At each timestep, the delta value represents the difference in measured current at that timestep and the measured current from the previous 10-second timestep. This acts to center the data about the x-axis and only consider deviations, e.g. an electric vehicle plugging in or a PV decreasing output as a string of thick clouds passes by. It also acts to place all feeders on the same playing field, e.g. a 15 A rise in demand on a lightly loaded feeder would be easier to notice than the same rise in demand on a heavily loaded feeder if the actual current values were used. The end of charging will not be able to be detected explicitly due to the ramping nature of the end of the charging cycle, but this fact should be useful for the BRNN to differentiate between charging events and non-charging events, e.g. a toaster or water cooker would have a high current spike when switched on followed by a corresponding negative current spike a few minutes later when it suddenly turns off. Fig. 3 and Fig. 4 show the 10-second delta values for one random month for Feeder #1 and Feeder #2 respectively.

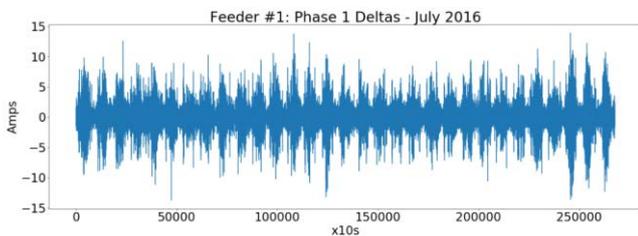


Fig. 3. Feeder #1: Phase 1 Delta Values - July 2016.

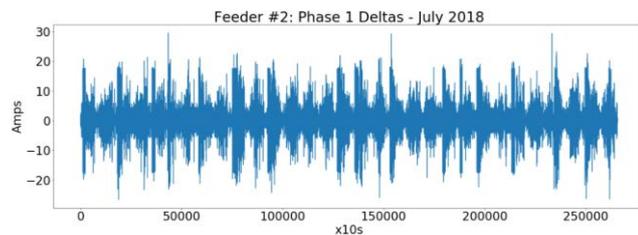


Fig. 4. Feeder #2: Phase 1 Delta Values - July 2018.

It can be seen that feeder #2 has much higher delta values than feeder #1. This difference in delta values cannot be clearly seen when considering actual current amplitudes. It can also be observed that for feeder #1, the model could simply fit a straight line at 15 A and would reach 100% recall and 100% precision. Precision and recall are the metrics used throughout this paper and are described in (1) and (2) respectively.

$$\text{precision} = \frac{\text{true positives}}{\text{true positives} + \text{false positives}} \quad (1)$$

$$\text{recall} = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}} \quad (2)$$

This is essentially the same for the rest of the data from feeder #1, therefore feeder #2 poses a much more difficult task and therefore will be the focus. Feeder #2 is a rural distribution feeder with a high PV penetration. As of 2016, there was a 400 kVA transformer supplying this feeder along with three other distribution feeders with a total of 189 households. There were 21 PV systems existing on the feeders with a total installed capacity of 160 kW. The phasor measurement unit (PMU) which was used to gather the data used in this experiment was intentionally connected to the feeder with the highest PV penetration out of the four feeders. The high penetration of PV on this feeder shows that residents connected to this feeder are open to new technology and energy solutions and therefore are also likely to purchase an EV in the coming years. Feeder #2 data was therefore used for training, optimization and testing purposes and feeder #1 was only used for transferability testing. The data was then converted into daily examples instead of monthly examples. The main driver behind using daily examples was to allow the DSO to view results every day instead of needing to wait an entire month, although there are also additional benefits in doing so. This is important for training because each example will be the same length and each timestep will be aligned. Although RNNs are capable of dealing with sequences of varying length, it makes more sense that each training example be a single 24-hour period where the model can learn time of day dependencies and learn what a single day should look like from start to finish than to learn an entire month. An entire month in 10-second resolution is also extremely long for a sequence, although it could be true that with the great power of the LSTM cell, this would still be possible. When using daily examples, each sequence consists of 8,640 timesteps, which is still very long and without LSTM cells would not be possible due to the vanishing gradient problem. The last step in the data preprocessing phase was to deal with any incomplete days and erroneous data points. After the data was finished being preprocessed, there were a total of 498 full days of three-phase data available for training, optimization, and testing, which corresponds to 1,494 examples after splitting up the three phases into separate sequences.

IV. METHODOLOGY VERIFICATION

Some EV charging data was collected from a charging station located in the parking lot of an office complex in order to see what the charging looked like after converting to 10-second delta values. The collected data proved that the

technique for transforming the ramp-up to full charging power into a single current spike was a feasible solution, although not perfect as the charging event could be spread out between two 10-second timesteps depending on when the vehicle was plugged in. Thus, lowering a given charging spike by as much as one half. Fig. 5 shows the effect on the delta values when an EV begins charging at various times within a 10-second interval. This was simulated by taking a single, real-life EV charging event and calculating the delta values ten times, each beginning the delta calculation with a one-second offset from the previous time. This shows that the delta value can vary depending on when charging begins between two consecutive timesteps. When starting the delta calculation at the third second, the delta value was reduced to a little greater than one half of the full charging power. However, at all other starting points, the delta value remained relatively high. This could likely still cause some false negatives to occur when deployed in a real-world scenario.

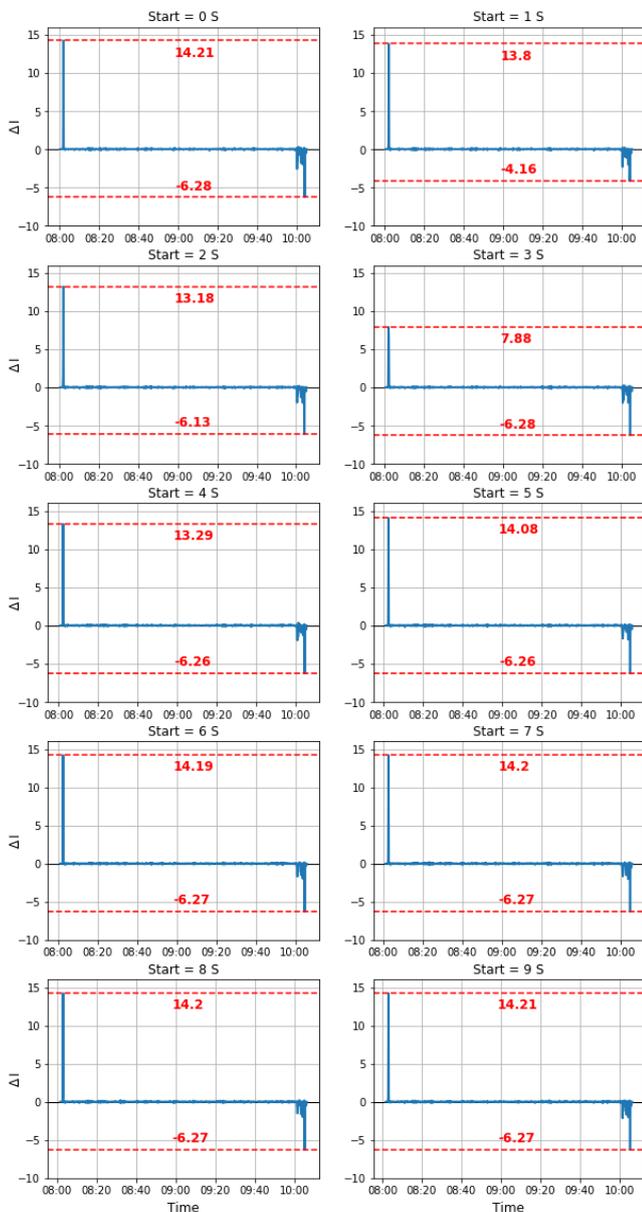


Fig. 5. Effect of Charging Start Time on Delta Value.

An issue caused by controlled charging could arise in the future if more charging infrastructure begins to implement charging methods which step the power consumption up and down as this could eliminate the “current spike” which is created by the fast ramp-up to full power. Other intelligent charging methods could also cause issues, e.g. if EV charging infrastructure is intelligently coupled to a rooftop PV system and the EV is switching between charging from the utility grid and the private PV system, as it could appear as multiple EV charging events. DG units connected to the distribution system could also cause similar issues if they are either fully or partially supplying power to EVs. This would eliminate or reduce the charging current seen by the measuring equipment at the secondary substation for a given feeder. Therefore, in the future, the structure of the injected charging events may need to be changed and made more complex. Not detecting EV charging which is supplied by a private PV system, for example, would not be a huge problem because it is not loading the distribution system significantly, but if the EV charging is supplied by some other low-voltage-connected DG unit then increased current levels would still exist on one or more distribution line(s) which could lead to a problem if left undetected.

V. OPTIMIZATION STRATEGY

For optimization of the model, the strategy was to create a model which would perform as well as possible on the training set first, i.e. low bias. Bias and variance are characteristics of the effective performance of the machine learning algorithm and are influenced by the model architecture, hyper-parameter tuning and the amount of data used to train the model. Analysis of bias and variance is essential to achieve high performance as it provides crucial information about how to optimize the model. Fig. 6 illustrates both high bias and high variance, where m is the amount of training data.



Fig. 6. bias and variance.

Refer to [11] for a more detailed explanation of bias and variance and how they affect performance. If the training set is not learned well, then the model will also not perform well on new data. Many of the hyper-parameters were optimized during this process aside from the number of hidden layers, the number of hidden units, the dropout rate (dropout was the only regularization method considered), and the prediction threshold. [12] is a great resource for understanding what a neural network is and to gain intuition as to how they function. Regularization is a widely used method of decreasing variance. Hyper-parameters, regularization and optimization of a neural network are explained in [11]. The remaining hyper-parameters were chosen temporarily as to minimize the loss on the training set. These hyper-parameters were then subsequently optimized to minimize loss on the

cross-validation set to decrease variance as much as possible while keeping the effect on the bias level to a minimum.

VI. OPTIMIZATION RESULTS

The best performing model was taken as the optimized model, which is laid out in Table 1. A slightly deeper model with a low dropout rate was found to perform marginally better when tuned correctly, but the recommendation given here would be to use the shallower network offered in this paper, as the requirements for computational power, memory capacity, and training time are significantly lower. Table 1 shows the hyper-parameters of the final optimized model.

TABLE I. OPTIMIZED NETWORK HYPER-PARAMETERS.

Affected Model Section	Hyper-Parameter	Result
Data	Amount of training data	896 samples (60% of maximum available)
	Number of features	1 (delta values)
	Feature scaling	Batch normalization
	Degree of pre-filtering	None
Model Architecture	Number of hidden layers	3
	Number of hidden units	4 (for all 3 layers)
	Hidden activation function	Sigmoid
	Output activation function	Sigmoid
Regularization (Dropout)	Dropout rate	0.0
Optimizer (Adam)	Beta 1 (normally not changed from default)	0.9
	Beta 2 (normally not changed from default)	0.999
	Learning rate	0.05/0.0001/0.00005 ^a
	Learning rate decay	0.0
Training	Number of epochs	700/70/20 ^b
	Batch size	128
	Loss function	Binary cross-entropy
Prediction	Prediction threshold	0.6

^a Trained with three different learning rates consecutively, starting with 0.05.

^b 700 epochs at a learning rate of 0.05, 70 epochs at a learning rate of 0.0001, etc.

As can be seen in Table 2, the final model was able to achieve a precision of 94.3% and a recall of 85.5% on the test set. In other words, the model was able to detect 85.5% of the EV charging events existing in the test set with a false positive rate of only 5.7%. Fig. 7 shows the block diagram for the final network.

TABLE II. MODEL PERFORMANCE ON TEST SET.

Test Set	
Precision (%)	Recall (%)
94.3	85.5

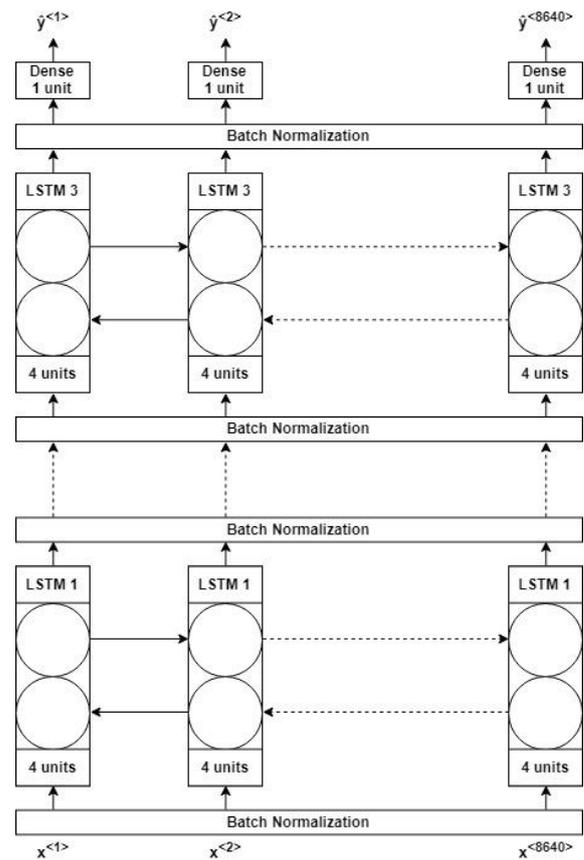


Fig. 7. Block Diagram of Final Model.

Fig. 8 and Fig. 9 show the false positives and false negatives per hour for the training set respectively. The model performed very well on the training data as expected, although there were some issues around 12:00 and 20:00 in both figures. This could have been caused by abnormalities in the data, as some data abnormalities were discovered at both 12:00 and 20:00 when looking at the average (average of all examples) delta values for each timestep. Fig. 10 and Fig. 11 show the hourly recall and hourly precision for the test set respectively, again showing reduced performance during midday. Fig. 12 shows the recall per EV charging amplitude using the test set. The bins with 0% recall signify charging amplitudes that simply did not exist in the data. The model, as expected, was able to detect higher charging amplitudes more easily.

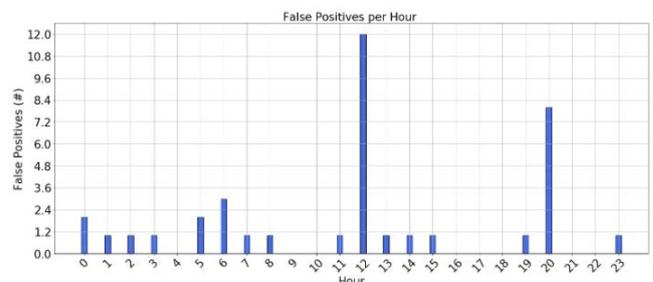


Fig. 8. False Positives per Hour - Training Set.

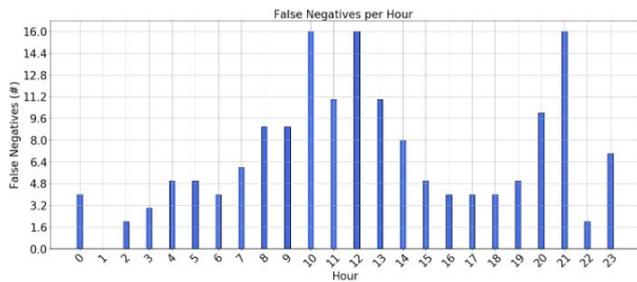


Fig. 9. False Negatives per Hour - Training Set.

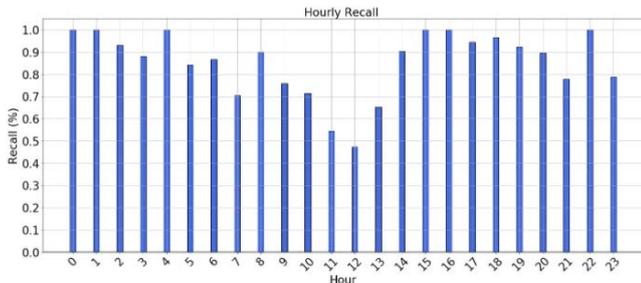


Fig. 10. Hourly Recall - Test Set.

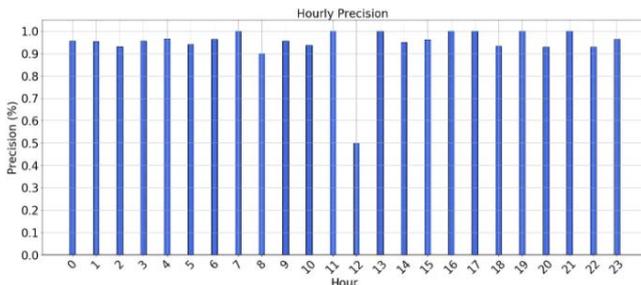


Fig. 11. Hourly Precision -Test Set.

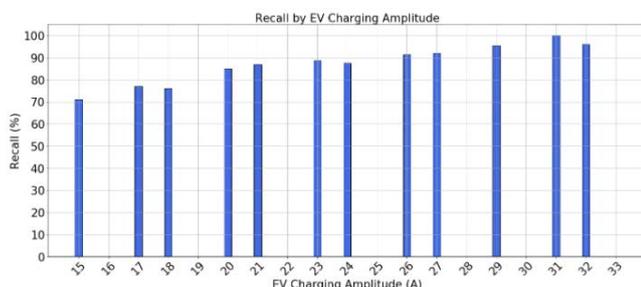


Fig. 12. Recall by EV Charging Amplitude - Test Set.

VII. TRANSFERABILITY

For this model to be useful in real life it is important to be able to train the model and have it work on different feeders. Otherwise, each feeder would need a substantial amount of measurement data before the model would be useful for detecting EV charging on that feeder and any change to the power system which changes the behavior of the feeder would require the model to relearn the new feeder behavior which would mean months of collecting data to retrain the model before being available for detection again. Unfortunately, there were only data for one other feeder available and this feeder was particularly easy to detect EV charging on due to the relatively small delta values existing on the feeder. Nonetheless, it was still another feeder for the model to be tested on.

Table 3 shows the performance of the model when evaluated on Feeder #1. Feeder #1 data was injected with EV charging events following the same procedure as Feeder #2. It shows that the model could perform very well when transferred to Feeder #1.

TABLE III. MODEL PERFORMANCE ON FEEDER #1.

Feeder #1	
Precision (%)	Recall (%)
100.0	98.6

During the development of this detection solution there was also a switching event which occurred, altering the topology of Feeder #2. The topology change provided another month of data for testing with different characteristics than the data used to train the model. The data was then prepared the same way as before except this time three different cases were made. In Case #1, the EV charging events were injected following the same procedure as before. In Case #2, a random number of charging events between one and three were injected into each day instead of a fixed two. In Case #3, only charging events in the range of 12 to 16 A were injected instead of 15 to 33 A. Table 4 shows the results for Case #1. The model performed well.

TABLE IV. MODEL PERFORMANCE ON TEST CASE #1.

Case #1	
Precision (%)	Recall (%)
100.0	98.3

Table 5 shows the results for Case #2. The results are very similar to Case #1, showing that the number of charging events does not affect performance in any significant way.

TABLE V. MODEL PERFORMANCE ON TEST CASE #2.

Case #2	
Precision (%)	Recall (%)
100.0	98.2

Table 6 shows the results for Case #3. The results here show that the model had a more difficult time detecting EV charging events of lower amplitude as expected. If using the model to detect low charging amplitudes, using a lower confidence threshold while predicting appears to greatly improve the detection rate while only decreasing the precision by a few percentage points. If the model is trained on lower charging amplitudes, it is also probable that it could perform better on charging events in this range, but even when trained exclusively on events in the range of 15 to 33 A, the model was able to detect events of lesser amplitude relatively well.

TABLE VI. MODEL PERFORMANCE ON TEST CASE #3.

Confidence Threshold	Case #3	
	Precision (%)	Recall (%)
60%	97.8	74.6
20%	94.3	84.8

Although there was not a large amount of data from different feeders, the above transferability tests show there is a high chance that the model is capable of being transferred to other feeders and performing well. In the future, when more data is available from various feeders with varying load profiles, it may be possible to improve performance. For example, by either training on multiple feeders at once to reduce the possibility of overfitting to one feeder or by saving many different models which have been trained on various feeder types so that a new feeder can be matched with a model which has been trained on a feeder with a similar load profile. Fig. 13 shows the actual number of charging events per charging amplitude which were injected into Case #3.

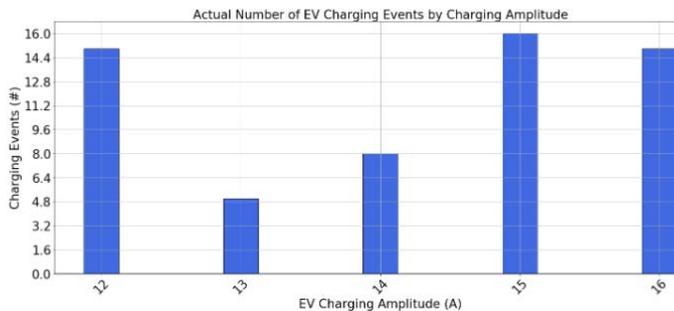


Fig. 13. Number of EV Charging Events per Charging Amplitude - Test Case #3.

Fig. 14 shows the performance of the model in relation to effective delta values. The amplitude bins indicate the effective delta values after the delta caused by EV charging is added to the existing delta at that timestep. This existing value could be a negative value or a positive value. In this case, there were no EV charging events that were reduced below one amp due to an existing delta value. It is not only the amplitude of the EV charging event at play but also the size of the existing delta value relative to the amplitude of the EV charging event. Lower amplitude charging events are simply more likely to be in the same amplitude range as existing delta values. Looking at Fig. 15, which shows the recall of the model in relation to charging event amplitude, it can be seen that there was less of a positive correlation between increasing charging amplitude and recall than there existed between effective delta value and recall as was seen in Fig. 14.

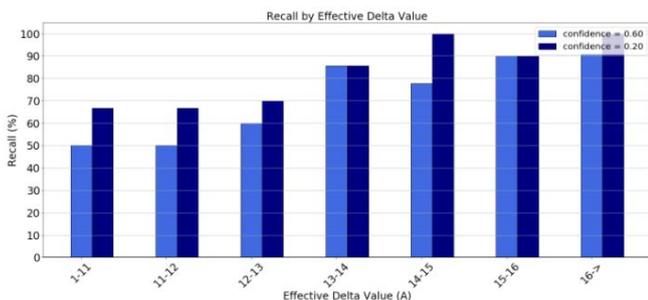


Fig. 14. Recall by Effective Delta Value, Test Case #3.

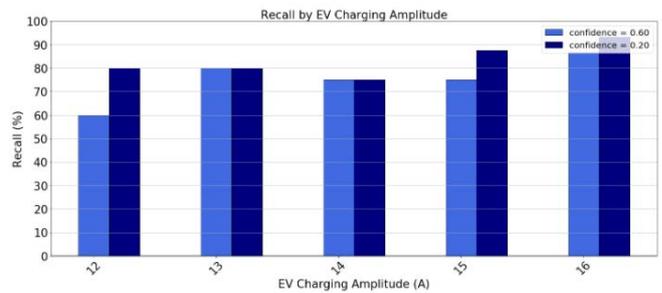


Fig. 15. Recall by EV Charging Amplitude, Test Case #3.

VIII. CONCLUSION AND FUTURE DEVELOPMENT

Through this paper, a machine learning model was successfully developed to detect charging of electric vehicles on a distribution feeder using only current measurements taken from the feeder at the low-voltage substation with high performance. The model can detect approximately 85% of the EV charging events in the test set (with charging events in the range of 15 to 33 A) along with some false positives. False positives accounted for 5.7% of the total number of predicted charging events, meaning around 5.7% of the detected charging events were not actually charging events. The model has also exhibited robustness against switching events as the model was able to detect EV charging events using test data created from the same feeder as before, but after an actual switching event took place. This is also an initial sign of transferability to other feeders, which was then reinforced by the positive results obtained when testing the model on data from a different feeder. Contrary to the initial hypothesis, the model was also capable of detecting EV charging events below 15 A with only a slightly diminished detection rate. The goal of the experiment was successfully achieved, but there are still many future development needs. The lack of measurement data from various feeders for transferability purposes along with the uncertainty of how the model will perform when evaluated on a feeder with actual EV charging events are probably the two biggest hurdles. Therefore, the first future development steps should be gathering more data and carrying out pilot projects where one or more known EVs are periodically plugged in at recorded times so that the model can be evaluated on its ability to detect actual EV charging events. These developments are required before a final product can be obtained. Continual developments will also be needed even after a final product is obtained to keep the product current with the latest technology, EV trends and changes in laws and other regulations that could affect the distribution system or EV charging behavior. Smart meter communication, changes in EV charging profiles due to controlled charging requirements and incentives, battery technology, consumer charging habits and other EV trends such as trends in charging amplitude and charging infrastructure are all prime examples of issues to monitor and stay up-to-date on in order to keep the product relevant in the future as well as the present.

ACKNOWLEDGMENT

I would like to thank Dr.-Ing. Eckehard Tröster, CEO and Senior Consultant at Energynautics GmbH, for giving me the opportunity and facilities required to carry out the necessary work for this paper and for his advice throughout the project. I would also like to thank Sabrina Hempel, Physicist at Energynautics GmbH, for her technical support throughout the project.

REFERENCES

- [1] "Was ist die Vision für die Zukunft? Aus welchen Bausteinen besteht das System Elektromobilität?", nationale-plattform-elektromobilitaet.de. [Online]. Available: <http://nationale-plattform-elektromobilitaet.de/hintergrund/>. [Accessed: 22- Mar- 2019].
- [2] "Moderne Messeinrichtungen / Intelligente Messsysteme", bundesnetzagentur.de. [Online]. Available: <https://www.bundesnetzagentur.de/DE/Sachgebiete/ElektrizitaetundG> as/Verbraucher/NetzanschlussUndMessung/SmartMetering/SmartMeter_node.html. [Accessed: 25- Mar- 2019].
- [3] A. Herzig et al., Smart Grid 2016, Die Digitalisierung der Energiewende. Deloitte, 2016.
- [4] A. Ng, "Machine Learning", Coursera, 2012.
- [5] A. Ng, "Sequence Models", Coursera.
- [6] Z. Lipton, J. Berkowitz and C. Elkan, A critical review of recurrent neural networks for sequence learning. University of California, San Diego, 2015.
- [7] R. Pascanu, T. Mikolov and Y. Bengio, On the difficulty of training recurrent neural networks. 2013.
- [8] W. Koehrsen, "Beyond Accuracy: Precision and Recall", Towards Data Science, 2018. [Online]. Available: <https://towardsdatascience.com/beyond-accuracy-precision-and-recall-3da06bea9f6c>. [Accessed: 08- May- 2019].
- [9] A. Ng, "Structuring Machine Learning Projects", Coursera.
- [10] T. Thomas and P. Plumhoff, "Power Quality in der Elektromobilität", Fachhochschule Bingen, 2013.
- [11] A. Ng, "Improving Deep Neural Networks: Hyperparameter Tuning, Regularization, and Optimization", Coursera.
- [12] A. Ng, "Neural Networks and Deep Learning", Coursera.