Explainable Anomaly Detection in Renewable Energy Power Plants by Learning Multidimensional Normality Models

Carsten Kleiner[1,*,†]

[1]University of Applied Sciences &Arts Hannover, Faculty IV, Ricklinger Stadtweg 120, 30459 Hannover, Germany

Abstract

Renewable energy production is one of the strongest rising markets and further extreme growth can be anticipated due to desire of increased sustainability in many parts of the world. With the rising adoption of renewable power production, such facilities are increasingly attractive targets for cyber attacks. At the same time higher requirements on a reliable production are raised. In this paper we propose a concept that improves monitoring of renewable power plants by detecting anomalous behavior. The system does not only detect an anomaly, it also provides reasoning for the anomaly based on a specific mathematical model of the expected behavior by giving detailed information about various influential factors causing the alert. The set of influential factors can be configured into the system before learning normal behaviour. The concept is based on multidimensional analysis and has been implemented and successfully evaluated on actual data from different providers of wind power plants.

Keywords

Anomaly detection, Attack detection, Resiliency, Multidimensional analysis, Wind power plant, Normality model, Explainable anomaly detection

1. Introduction and Motivation

For reasons of sustainability the amount of regenerative power production is continuously increasing worldwide at ever higher rates. With higher shares of the overall power production, the importance of a reliable power supply from renewable sources becomes more and more important. On the other hand, due to their dependence on actual weather conditions, it is more difficult to achieve a reliable supply from natural sources as a matter of principle. Thus, an even closer monitoring of the production process by the operators is important to account for that.

Apart from operational challenges, the rising impact of renewable sources in power production also makes them an attractive target for attackers to achieve evil purposes. As already shown by the attack on Ukrainian power plants in December 2015 by Russian hacker groups, critical infrastructure becomes an ever more important attack target, not only in the recent war crisis in Ukraine ([1]). Thus, it is also important to employ advanced and powerful attack detection systems for renewable power production systems in order to protect this part of the critical infrastructure.

In this paper a novel detection system will be proposed that is capable of detecting anomalies in the operation of renewable power plants. The system operates reason-agnostic in its ability to detect anomalous operation, be it operational or based on attacks. Since monitoring and decisions on potential actions to be taken are ultimately performed by highly skilled humans, it is important to use their time as economically as possible. By integrating outage and attack detection in a single system, this goal is supported.

In addition, typically there is a tradeoff between false positives and false negatives to be balanced in anomaly detection. The more alerts are generated, the smaller the number of false negatives. On the other hand, more alerts often means more false positives, exhausting the human resources to deal with the generated alerts. Thus, in order to take informed decisions and apply appropriate measures, the human monitoring staff needs to be able to assess messages from the anomaly detection engine. So it is important that reasons for alerts are provided to the humans in order to detect false positives as easy as possible. The proposed system will provide such reasons to the operators by showing detailed, mathematically based explanations for generating alerts.

The remainder of this paper starts with a review of related publications in section 2 which will show that while there are already advanced solutions to specific aspects, none of these systems provides the combination of features as our system. The concept of the proposed system will then be explained in section 3 and specific configuration for wind power plants will be presented. This is followed by a practical evaluation of the concept on actual wind power plant data from different German wind power plants from years 2019 to 2021 in section 4. Finally, results will be summarized and ideas for extending the system itself as well as its application scope will be presented in section 5.

Published in the Proceedings of the Workshops of the EDBT/ICDT 2024 Joint Conference (March 25-28, 2024), Paestum, Italy

ckleiner@acm.org (C. Kleiner)

0000-0001-9497-0312 (C. Kleiner)

2. Related Work

Several papers in the context of anomaly detection for renewable energy systems can be found in the literature. In a more generalized context, [2] describes a learning approach similar to the one in this paper even for any type of IoT system. Whereas this approach could also be applied to renewable power plants, it is not clear which part of the learning can be carried out in an automated fashion. Similarly, results do not provide explanations for anomalies. A focus on attacks, more specifically intrusion detection, is described in [3]. However, the approach is not extensible to outage detection and only provides non-explainable alert messages. More specifically for power plants, [4] uses many very general input parameters. However, this approach also does not provide explainable anomalies as results.

Other interesting wind power specific concepts include [5, 6]. However, these approaches also do not provide explainable results. The first, in addition, requires a semi-supervised learning approach which is not feasible for previously unknown attack types. Also, annotated training data is often times not available. The second approach focuses on system failure detection rather than attacks.

On the other hand, [7] focuses on attacks and is specific for wind power plants. It is not extensible to other types of energy sources and the degree of explainability of the results is not obvious. Papers [8, 9] also only focus on specific attacks for wind power plants and thus do not achieve the general detection capabilities of our concept. The latter is concerned with false data injection attacks which are also the focus of several other publications. Moreover, [10] provides a good overview of the security challenges from attacks that have to be considered, but it does not present a comprehensive solution.

Finally, there are also papers with a pretty similar concept to ours, but with different detection approaches, such as Markov chains in [11] and a more complex detection model in [12]. However in both cases, while the approach is specific to wind power plants and an extensibility is not documented, the explainability of the generated alerts is uncertain. This is also true for [13] which also uses a correlation based approach, yet it is only one-dimensional and requires and includes many specific sensors, so that it is also tied to the domain of wind turbines only. Even more specific to wind turbine gearboxes is [14]. The authors do not limit their approach to attacks, also use a multidimensional analysis and generate at least partially explainable alerts. However, it is not obvious whether and how this can be extended beyond gearboxes.

In summary, none of the discussed references is able to provide the comprehensive features of our approach (cover attacks and outages, generate explainable alerts, capable of detecting unknown attacks and useable for different types of power generation).

3. Concept

3.1. Requirements and Context

Based on the research project SecDER[1] which aims to increase the resilience of renewable virtual and physical power plants, the requirements for an anomaly detection system have been identified as follows:

Reason agnostic Both anomalies originating from known and unknown attacks as well as non-attack based anomalies shall be detected, ideally based on a single detection system.

Explainable alerts The identified anomalies should be used to raise alerts that can be handled by human domain experts. In order to simplify and substantiate the decisions by the experts explainable alerts should be provided, detailing the reason and context why the alert has been issued.

Adaptability The concept shall be usable for different types of wind power plants as well as different types of renewable power plants in general. The learned normality models can be specific for each plant, however, the concept to learn the model should be generic.

General normality model While a single set of normality models for all plants is not a goal, it is preferable, if normality models can be learned for groups of similar plants. This way the model becomes more stable, and the number of extensive learning processes can be reduced.

Continuous learning and adjustment The system should be capable of adjusting the learned system behaviour continuously, thus improving the quality of the normality models over time. Thus can also update the models in cases of concept drift over time.

The system described in the following part of the paper will satisfy all of these requirements. On the other hand, there are also limitations of the approach that have been accepted in order to keep the complexity manageable. In particular, detection is only considered up to explainable alert generation, alert handling itself is not in scope. Handling can be considered orthogonal as long as explainability of the generated alerts is secured. For alert handling, generic procedures and manual update concepts can be considered as an extension, see e. g. [15] for an approach based on rule-based anomaly detection. Similarly, we only consider anomaly-based detection concepts, since most attack patterns (and even some of the non-attack-based outage patterns) are previously unknown, so rule- or pattern-based detection will not be powerful enough to detect these. As attacks on virtual power plants are executed by designated experts, advanced attacks will be used which are unique to the specific target and thus typically not previously known.

[1] https://secder-project.de

3.2. Multidimensional Normality Models (MNM)

The basic concept for anomaly detection is learning multidimensional normality models (MNM) based on historic data of the power plant (or a set of similar power plants) and then assessing the deviation from this MNM for current readings of a logical record of the plant. The concept called cellwise estimator (CE) of the MNM has already been described in [16] in detail; thus, we will only present a high level description here. Originating from online analytical processing (OLAP) cubes, the idea is to describe normal behaviour of certain metrics (such as power production in a windmill) based on several orthogonal dimensions (such as weather conditions, plant sensor readings and others). The reason for this multidimensional treatment is that measurements of the metrics may be within a permissible range when looking at them globally, whereas they may be an anomaly, when considering the specific context in more detail. The context is described by the dimensions which are used in learning the MNMs. Conversely, potentially abnormal measurements on the global level may actually be normal when looking at their specific context. Thus, it is important to be able to base a decision whether a logical record constitutes an anomaly on both global as well as contextual, i. e. dimensional, information. To account for these challenges a specific normality model is learned for each of the cube cells, i. e. every contextual situation.

Unfortunately, the higher the number of dimensions and the number of values within a dimension, the larger the number of combinations to consider becomes. Since the growth is exponential, these numbers have to be limited. In addition the concept of iceberg cubes ([17]) known from the OLAP domain can also be used to restrict the number of cubes to consider to relevant ones.

In order to deal with continuous data streams as needed for monitoring a power plant, the cubes are computed per timeslice with a configurable timeslice length. The metric attribute whose normal behavior is to be learned is aggregated by some configurable aggregation function over all readings within a timeslice. For the domain of wind power plants for instance, the power production output of a mill is a logical choice as a metric with multiple readings being aggregated by using the average over a timeslice. Typical dimensions for this metric can be wind speed, wind direction, rotor position and outside temperature. Since the dimensions are used to form an OLAP-like cube, all dimensions must be of discrete types. Thus, continuous readings such as wind speed and temperature need to be assigned to a set of classes in order to be used as dimensions. As known from OLAP rollups, there is also a symbolic value of * in each dimension that aggregates all classes in that dimension and thus provides a cube cell where the class is irrelevant.

The goal of the learning process by looking at historical data is to compute a statistical description of the metric attribute for each cell of the cube. This is done by assuming a normal distribution for the metric readings in each cell and approximating that normal distribution by estimating mean and standard deviation for the metric attribute based on learning from historical data. For current readings the anomaly score is computed as difference to the mean of each relevant cell as number of standard deviations. The higher this factor, the more likely the current reading is an outlier. As known from statistics a factor of 3 is a natural choice as a threshold to generate an alert. As will be seen in section 4, solely looking at this factor as an anomaly measure is not sufficient, though, to properly assess the importance of an alert.

In summary, each cell's normality model in our concept consists of an estimation of normal distributions (with mean and standard deviation each) of one or more measurements per cube cell over a timeslice. Cube cells are defined by combinations of discrete values of relevant dimensions, with wildcards allowed for cells with irrelevant values in a dimension. The anomaly score is then computed based on the number of standard deviations that any current reading of a measure deviates from the expected mean. Alerts are typically only raised for cube cells with anomaly scores higher than a threshold of 3. In addition to the anomaly score the computed normality model as distribution estimation is also provided with the alert along with information about the cell's dimensional values that caused the alert. This combination of information (metric measurement, anomaly score, contextual values, normality model) comprises the explanation for the human expert. Thus, an informed decision about proper reaction to the alert is facilitated.

3.3. Application of MNM to Wind Power Plants

In order to apply our concept as explained in section 3.2 to renewable energy plants in general and wind power plants in particular, we have to define the metrics with aggregation functions for which normality models shall be learned as well as the discrete influential dimensions that might influence the metrics and be important for assessing an alert. Candidates for choosing the metrics are any elements of a monitoring reading that can be used to describe the operational behaviour of a windmill. The assumption is that attacks or outages will lead to unexpected behavior in this metric. Primarily, this is the effective electrical power production of the mill computed as an average over a timeslice. For consistency checks the number of measurement readings per timeslice can also be used as a metric. Alternative options that have not been evaluated in the experiments described in

section 4 could be the positions of the pod or the blades of the windmill or other operational features.

There are much more options for choosing the dimensions than the metrics. In the evaluation in section 4 we have experimented with different choices, but there are actually many more. Obvious dimensions include wind speed, wind direction, pod position, air temperature, air pressure. More possible options include power factor, pitch angles of each blade, angle between pod and wind direction and anemometer readings. The choice of discretization of each of these factors (cf. 3.2) can be considered another hyperparameter of the application. Specific choices for the dimensions and discretizations for the experiments will be explained in section 4, but it has to be pointed out that those are only initial selections and much more experiments will have to be carried out in the future to optimize the approach, cf. section 5.2.

4. Evaluation

In order to evaluate the capabilities of the concept in detail, we used historical data from actual wind power plants that are operated by project partners in the SecDER project. We had two different datasets, one from each operator. Data did not contain any known attacks, yet some anomalies due to maintenance or unusual weather conditions.

The first dataset consists of operational log data from a single wind mill over the time range from January 2020 to August 2021 at a sampling rate of 15 minutes. Each log reading consists of 22 attributes in total, one of which is the timestamp and the others can be used as metrics or dimensions as will be explained in section 4.1.

The second dataset provides operational log data from 9 different wind parks, comprising 42 windmills in total at a sampling rate of 5 minutes. Data provides 30 attributes per reading and readings were available for the year 2020.

In both cases, a first part of the data has been used for training and the remainder for testing. In the sequel, results will be presented based on output from a specifically developed GUI tool. In the figures the testing period will be used horizontally to display the results for individual test instances. Each timeslice's reading can be considered a test case. The graph shows the results for a specific cell of our cube, as selected from different dimensions, values and combinations at the top. Within a figure the red curve shows the computed metric value (scale on left) whereas the blue curve shows the anomaly score (i. e. the number of standard deviations that the value is from the mean in this particular cell), scale on the right. Typically, scores above 3 can be considered anomalous. In addition, a yellow line displays the learned mean value for the metric for this cell and green and lightblue lines show mean +/- 3 standard deviations.

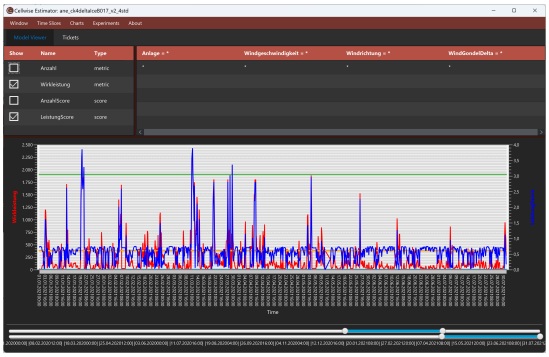

Figure 1: Effective Power and Anomaly Scores for Single Mill (Total view)

4.1. Validation of Concept

As an initial validation we used the data from 2020 of the first dataset as training set and the readings from 2021 for testing. We chose the average electrical power production over timeslices of 4 hours as primary metric. We experimented with some attributes as dimensions, the results in this subsection have been achieved with wind speed, wind direction and difference between gondola angle and wind direction. The continuous values in these dimensions have been linearly assigned to 9, 12 and 5 classes, respectively. The number of classes of the first two features has been determined heuristically by assigning equally sized intervals of the total range of values to classes. For the third feature where original data had a strongly non-linear distribution we decided to use fewer classes to primarily account for major and medium outliers in each of the two directions and have most data in the no difference class.

Figure 1 shows the test results for the global cell, i. e. no fixed value in any of the dimensions. As we can see, there are only few significant anomaly scores, primarily those on January 20th, March 11th and March 29th. At this general level (no fixed dimensional values), this behavior can be expected as the threshold for raising an alert is around 1900 kW which is already pretty close to the 2400 kW nominal power of the mill. However, the first two of those scores will not be reported by an alert as all subcells into the wind speed direction do not have an anomalous score. This means that the power production seemed unusually high from a global point of view (which is information that could have been observed without our approach but would have raised a false positive), yet in reality it is simply explainable by the rather high wind speed on those days. For the remaining high anomaly score the dimensional analysis shows reduced anomaly scores the further detailed the cells become, yet it remains above 3, thus raising an alert. Looking at the data in

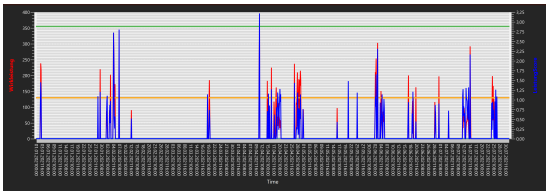

Figure 2: Effective Power and Anomaly Scores (Single mill, Two dimensions restricted view)

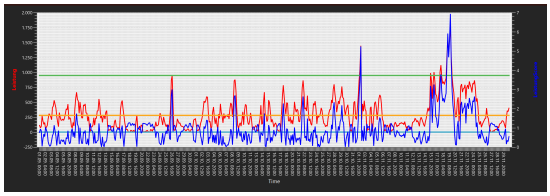

Figure 3: Effective Power and Anomaly Score (Plant Group, total view)

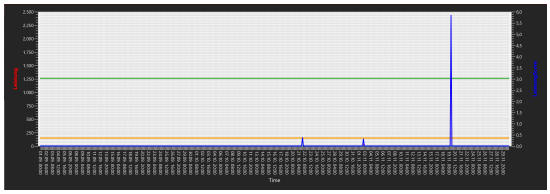

Figure 4: Effective Power and Anomaly Scores (Plant group, dimensionally restricted view, speed class 6, direction class 8)

detail in the evaluation, this score can be considered a false positive. The reason is that this specific context situation had not been observed in the whole training period. Such errors can be remedied by increasing the training data set.

Even more interesting is the analysis looking into some of the dimensions, as the learned normality behavior is much more specific in those cases as seen in figure 2. In that figure we have focused the display on the wind speed class 2 (pretty low speed) and the wind direction class 2. The figure shows that the learned model with mean around 140 kW and 80 kW standard deviation is very specific. Still, the only remaining alert with an anomaly score of 3.1 shows up at April 11th. This could be a false positive due to a too specific cell model or a true alert due to a malfunction with too high generated power. A human operator seeing the alert would be able to classify this alert based on his domain knowledge. Due to space constraints we only present these exemplary results here.

4.2. Common Model for Plant Groups

For the second validation data from the set of windparks has been used. Here, January to August 2020 has been used as training data and September to December 2020 for testing. Metrics and dimensions shown are identical to the ones in the previous subsection for comparability purposes. In addition, the specific wind mill has also been used as another dimension in order to be able to analyze the outcome per mill and over all mills together. Data from 17 of the mills with identical nominal power production of 2300 kW have been used.

Figure 3 again shows the overall view of the scores with no fixed dimensional values. We can see that the learned normality model is much more specific than the one in figure 1 due to the extended training set (standard deviation around 200 kW as opposed to 500 kW).

Two cases with higher anomaly scores can be identified, namely Nov 2nd and Nov 19th/20th. The first of those shows a similar behavior as already noted in the previous subsection, i. e. an anomaly score that does not show up in any of the dimensionally restricted models and thus, it would not be reported as alert. The latter anomaly score would be tied to two of the four windparks as well as specifc wind speed and wind direction all showing anomalouss scores in one alert as those are all dependent cells in the cube. This shows that the score is indeed an anomaly for these mills (cf. figure 4) and should thus be reported as an anomaly alert. This can be considered a true positive that is recognized by the system. It can be further explained to the human expert by providing the specific wind park, speed and direction that causes the alert to be raised.

In general, the increased size of the training data leads to more precisely learned models in the cells. This potentially increases the number of false positives, since anomaly scores are more likely with smaller standard deviation. However, by judging an anomaly score in combination with the standard deviation of its cell, most of the false positives can be identified easily and thus do not lead to raising alerts. On the other hand the benefit of the more precise models is that false negatives are much less likely in that case.

Also, only precise cell models facilitate discovery of anomalies in cases with unusual low power production particularly relevant in case of attacks. This is due to the fact that low production is only observed as an anomaly if the learned mean - 3 standard deviations is above 0 kW. This can only be achieved with rather precise cell models which need large training datasets.

4.3. Evaluation against Known Outages

The evaluations in the previous subsections were only able to show that anomalous behavior can be detected in principle, since the data did not contain any *known* attacks or outages of the power plants. In order to get a qualitative impression of how well the detected anoma-

lies correspond with actual unusual behavior, we evaluated the concept against data from a single windmill that was available over a 2.5 years time frame. In addition, for this plant information from the plant management system (PMS) was available that listed all known *and recorded* system problems during that time.

It should be noted that this evaluation is not well suited for a thorough quantitative analysis of the algorithm since the dataset only provides information about events affecting the operation of the mill that were known to the PMS. Thus, since no attacks are known there are no attack labels and thus no evaluation against attack detection is possible. Similarly, anomalous situations due to an unusual behavior of the mill unknown to the PMS are not labeled as anomalous in the ground truth. Thus we can expect some (seemingly) false positives for the anomalous situations not recorded in the PMS and thus labeled as normal. This will lead to a rather low precision when comparing our anomaly messages with the events recorded in the plant management system as ground truth.

In addition, the events in the PMS record any unusual situation in the windmill regardless of their impact on the actual power production. Since we consider output power production as our analysis target, it is obvious that we will not be able to detect events that have no or minimal influence on the power production[2]. Such situations will be recorded as seemingly false negatives in the comparison, impacting the recall negatively. However, we do not anticipate too many of such messages so that aiming for a high recall is still a desirable target.

Both effects mentioned previously will also impact other measures such as accuracy (to some degree) and F1 score (to large degree). Still a good, albeit not perfect, accuracy score is also a valid goal to target.

4.3.1. Evaluation Setup

For this evaluation we used windmill data from 2 years as training set for our algorithm and data from the remaining 0.5 years as a test set. We used an algorithm configuration similar to the one in section 4.1. We had to clean training data by removing the readings for times which had been recorded in the PMS as anomalous in order to only learn normal behavior of the system.

Since the events recorded in the PMS used timestamps with 5 minute difference, we first need to align the time resolution, i. e. define how many anomalous events within a 4 hour timeslice make such a timeslice anomalous in total. While it is desirable on one hand to even realize anomalies that only occur at a single instance in time, on the other hand it is questionable whether a full timeslice shall be considered anomalous just based on a single event. For the following evaluation we used thresholds of 40 and 5 minutes within a 4 hour timeslice as a condition for an anomalous timeslice. Note that an anomaly due to an outage is usually rarely a very short incident.

[2]From a practical point of view detecting such events with our algorithm is not necessary, as these have only minimal impact on the power production and are already known from the PMS and thus do not require advanced detection.

PMS issue	false	948	103
	true	15	38
CE anomaly alert		false	true

Table 1
Confusion matrix for outage anomaly detection (at least 40 minute outage per timeslice considered anomalous)

Another aspect is the management of missing readings from the windmill which is often times caused by anomalous operation. If no data readings are present for a whole timeslice the CE algorithm will not detect an anomaly for the power production, since missing data does not get any anomaly score. However, with the second metric (number of readings per timeslice) we can easily detect timeslices where no power readings are present and thus report them as an anomaly as well. Finally, a single anomalous cube cell per timeslice will make the entire timeslice anomalous. This is one of the primary strengths of the algorithm to also detect only specific anomalies within a large set of non-anomalously seeming other cells at the same time. The explaination of the anomaly for a timeslice will contain all anomalous cube cells for that timeslice together with the additional data, so that the human expert can further examine the incident.

4.3.2. Exemplary results

With the setup as described before and 40 minute anomaly threshold we achieved a recall of 0.72 and an accuracy of 0.89 as the primary targets of the algorithm. The precision was low at 0.27 as expected and explained above; this makes an F1 score of 0.39. The matrix in table 1 summarizes the results.

Again, the seemingly high number of false positives is due to the fact that the CE detects anomalies that are not part of the PMS failure ground truth, either because they are attacks or because they did not lead to events in the PMS. As another baseline an auto-encoder based algorithm trying to detect only outages on the same data set only achieved a 0.31 F1 score, mainly because of a higher number of false negatives.

If we reduce the threshold how many anomalous events in the groud truth make a timeslice anomalous to a single event (i. e. 5 minutes of the 4 hour timeslice), the recall reduces somewhat to 0.60, however accuracy and precision remain pretty much the same such that the F1

PMS issue	false	938	103
	true	25	38
CE anomaly alert		false	true

Table 2
Confusion matrix for outage anomaly detection (at least 5 minute outage per timeslice considered anomalous)

PMS issue	false	1004	47
	true	17	36
CE anomaly alert		false	true

Table 3
Confusion matrix for outage anomaly detection with higher anomaly threshold

score reduces to 0.37 (cf. table 2). This behavior is due to an increased number of false negatives, which could be expected as some minor issues in plant operation do not necessarily cause anomalous power production. The auto-encoder baseline increased its F1 score to 0.33 in this case.

A final evaluation shows that there is still potential in the CE based algorithm by fine tuning the learned cell models. Increasing the threshold anomaly score for alerts to 4 standard deviations, we obtain the confusion matrix in table 3. This increases the accuracy to 0.94 and specifically the precision to 0.43. The recall is slightly reduced to 0.68 for an overall F1 score of 0.53. This improvement is primarily due to the reduced number of seemingly false positives in situations where no outage is recorded in the PMS. However, it remains unclear whether this is an actual improvement in practice or not. It simply leads to a reduction of detected anomaly candidates. Yet from the data provided it is unknown where these situations would actually belong to anomalous or regular behavior.

In summary, the evaluation in this section has shown that the algorithm introduced in chapter 3 is capable of detecting unusual system behavior of a wind power plant which had also been recorded in a PMS, particularly with good accuracy and recall. Precision and thus F1 score are somewhat lower which can be attributed to the algorithm also detecting anomalous behavior that had not been recorded in the PMS, e. g. because it was due to a specific wind condition. This is exactly what the main advantage of the CE algorithm is, namely also detecting anomalous behavior in specific conditions which could be caused by an attack. We have also shown optimizing some of the hyper parameters of the approach (such as message thresholds and timeslice aggregation) might improve the detection quality further in addition to larger training sets and more dimensions.

5. Conclusion and Future Work

5.1. Summary

In this paper we have presented a concept and implementation to detect anomalous behavior in renewable power plants. The concept is based on learning normal behavior of key performance figures such as effective power production. The normal behavior is learned for many specific situations which can be expressed as multidimensional cells in an OLAP-like data cube. On one hand, this reduces the number of false negatives by learning very specific models for the individual cells representing specific situations. On the other hand, the number of false positives can still be kept low by using larger training data sets. Also, assessing the specificity of the learned model to put a mere anomaly score into context and thus facilitate appropriate treatment before raising alerts can be done by a human inspector and to some degree even an automation such as in section 4.3. This is an important advantage of the explainability achieved by the learned behavior models for each cell. The concept has been successfully evaluated on actual data from wind power plants as shown in section 4 both in general and also on a set of known outages as one possible reason for anomalous behavior.

In summary, the concept presented in this paper offers a promising approach to detect anomalous behaviour in renewable power plants by learning specific models according to a configurable set of dimensions reflecting relevant circumstances for power production. The anomaly scores based on learned mathematical models provide traceable explanations for the detected anomalies which may originate from attacks or regular operational issues.

5.2. Outlook

While the evaluation presented in section 4 already showed the usefulness of the concept, much more experiments are needed to reveal its full potential. Much more analysis with regard to identifying interesting and relevant dimensions in the base data to be used for the cube is required. Some promising dimensions such as temperature, air pressure and power factor have not been included yet. Moreover, using larger time ranges for the training data will be one of the next steps to further verify the positive impact of more precisely learned models. This should also further reduce some issues detecting unusual low power production due to normality models with too large standard deviations that do not raise high enough anomaly scores even for zero power production in certain situations.

Also, some experiments have shown that using a normal distribution as foundation of estimating cell models is not always appropriate. We saw several cases where

most metric training data lies around a rather small value with a few high outliers. For such distributions a normal distribution is not a good estimator. Instead, alternative models should be used which will be added to our implementation soon.

Finally, we have currently only evaluated the concept on wind power production. We have similar datasets from photovoltaics which we plan to use for a second evaluation. Metric will be similarly the effective power production, but regarding dimensions there will have to be an extensive evaluation which are most promising.

References

[1] C. . I. S. Agency, Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors, 2018. URL: https://www.cisa.gov/uscert/ncas/alerts/TA18-074A.

[2] S. Chakraborty, A. Onuchowska, S. Samtani, W. Jank, B. Wolfram, Machine learning for automated industrial iot attack detection: An efficiency-complexity trade-off, ACM Trans. Manage. Inf. Syst. 12 (2021). URL: https://doi.org/10.1145/3460822.

[3] K. N. Junejo, J. Goh, Behaviour-based attack detection and classification in cyber physical systems using machine learning, in: Proc. of the 2nd ACM Int. Workshop on Cyber-Physical System Security, CPSS '16, ACM, New York, NY, USA, 2016, p. 34–43. URL: https://doi.org/10.1145/2899015.2899016.

[4] P. Sun, J. Li, Y. Yan, X. Lei, X. Zhang, Wind turbine anomaly detection using normal behavior models based on scada data, in: 2014 ICHVE International Conference on High Voltage Engineering and Application, 2014, pp. 1–4. URL: https://doi.org/10.1109/ICHVE.2014.7035504.

[5] Y. Zhou, W. Hu, Y. Min, et al., A semi-supervised anomaly detection method for wind farm power data preprocessing, in: 2017 IEEE Power & Energy Society General Meeting, 2017, pp. 1–5. URL: https://doi.org/10.1109/PESGM.2017.8273883.

[6] C. McKinnon, J. Carroll, A. McDonald, et al., Investigation of anomaly detection technique for wind turbine pitch systems, in: The 9th Renewable Power Generation Conference, 2021, pp. 277–282. URL: https://doi.org/10.1049/icp.2021.1401.

[7] H. Badihi, S. Jadidi, Z. Yu, Y. Zhang, N. Lu, Smart cyber-attack diagnosis and mitigation in a wind farm network operator, IEEE Transactions on Industrial Informatics (2022) 1–10. URL: https://doi.org/10.1109/TII.2022.3228686.

[8] A. Datta, M. A. Rahman, Cyber threat analysis framework for the wind energy based power system, in: Proc. of the 2017 Workshop on Cyber-Physical Systems Security and PrivaCy, CPS '17, ACM, New York, NY, USA, 2017, p. 81–92. URL: https://doi.org/10.1145/3140241.3140247.

[9] K. Guibene, N. Messai, et al., A data mining-based intrusion detection system for cyber physical power systems, in: Proc. of the 18th ACM Int. Symposium on QoS and Security for Wireless and Mobile Networks, ACM, New York, NY, USA, 2022, p. 55–62. URL: https://doi.org/10.1145/3551661.3561367.

[10] A. Jindal, A. K. Marnerides, A. Scott, D. Hutchison, Identifying security challenges in renewable energy systems: A wind turbine case study, in: Proc. of the 10th ACM Int. Conf. on Future Energy Systems, ACM, New York, NY, USA, 2019, p. 370–372. URL: https://doi.org/10.1145/3307772.3330154.

[11] J. D. Deng, H.-S. Lee, C. McMillan, A. Rimoni, M. Zhang, Analyzing wind speed data through markov chain based profiling and clustering, in: Proc. of the 2nd Workshop on Machine Learning for Sensory Data Analysis, MLSDA'14, ACM, New York, NY, USA, 2014, p. 67–73. URL: https://doi.org/10.1145/2689746.2689756.

[12] N. Song, X. Hu, N. Li, Anomaly detection of wind turbine generator based on temporal information, in: Proceedings of the 2019 7th Int. Conference on Information Technology: IoT and Smart City, ICIT '19, ACM, New York, NY, USA, 2020, p. 477–482. URL: https://doi.org/10.1145/3377170.3377271.

[13] H. Lee, N.-W. Kim, J.-G. Lee, B.-T. Lee, An approach for utilizing correlation among sensors for unsupervised anomaly detection of wind turbine system, in: 2021 Int. Conf. on Information and Communication Tech. Convergence, 2021, pp. 104–109. URL: https://doi.org/10.1109/ICTC52510.2021.9621198.

[14] S. Zhu, Z. Qian, B. Jing, M. Han, Z. Huang, F. Zhang, Condition monitoring of wind turbine gearbox using multidimensional hybrid outlier detection, in: Int. Conf. on Smart-Green Technology in Electrical and Inf. Systems, 2021, pp. 112–117. URL: https://doi.org/10.1109/ICSGTEIS53426.2021.9650387.

[15] L. Renners, F. Heine, C. Kleiner, G. Dreo-Rodosek, Concept and practical evaluation for adaptive and intelligible prioritization for network security incidents, International Journal on Cyber Situational Awareness 4 (2019) 99–127.

[16] F. Heine, Outlier detection in data streams using OLAP cubes, in: New Trends in Databases and Information Systems - ADBIS Short Papers and Workshops, Nicosia, Cyprus, volume 767 of *Communications in Computer and Information Science*, Springer, 2017, pp. 29–36. URL: https://doi.org/10.1007/978-3-319-67162-8_4.

[17] J. Han, J. Pei, G. Dong, K. Wang, Efficient computation of iceberg cubes with complex measures, SIGMOD Rec. 30 (2001) 1–12. URL: https://doi.org/10.1145/376284.375664.