Benchmarking a Physics-Based Approach for Anomaly Detection at Utility PV Plants

Scott Sheppard[1], Keith A. Dickey[2], Steven Koskey[1], Corson Teasley[1], Christopher Perullo[1], Daniel Fregosi[3], Wayne Li[3]

[1]Turbine Logic, Atlanta, GA, 30332, USA
[2]Southern Power Company, Birmingham, AL, 35243, USA
[3]Electric Power Research Institute, Charlotte, NC, 28262, USA

Abstract — Many utility monitoring and diagnostic centers have adopted advanced pattern recognition software to aid in anomaly detection and diagnosis. Due to the wide variety of electricity generation methods and associated supporting hardware, utilities choose software that is applicable to the broad category of industrial hardware. As a result, these tools excel at detecting large deviations from normal operation but struggle to identify subtle shifts in performance that are indicative of the onset of degradation and failure. At worst, these tools can be oversensitive and raise false alarms when the deviations are explained by operation outside of what was observed in the tool's training data. Recent developments in physics-based modeling have resulted in models that are capable of accurately detecting faults in the DC collector field that, individually, results in a less than 5% power loss at the combiner box level. These new models are benchmarked against current state-of-the-art utilities tools, with models designed to match the physics-based approach as much as is feasible. The applied physics-based models improve fault detection capabilities over the standard utility tool, detecting approximately twice as many real faults for a given false positive rate.

I. Introduction

The United States continues to see accelerated growth in installed PV generation capacity. Current EIA estimates predict that solar generation capacity additions in the U.S. will grow by 62% in 2024 [1]. Furthermore, the EIA also expects that almost all growth in U.S. electric generation through 2025 will come from solar power [2]. Under these projections, power generation from solar power is expected to exceed power generation from hydropower in 2024 [3]. While installed capacity continues to grow, current installations regularly underperform expectations. In 2021, both fixed-tilt and solar tracking systems fell short of projected power output by 8% [4].

One driver of this system-level underperformance is faults associated with a plant's DC hardware. While large-scale PV system hardware outages, such as inverter outages, are straight forward to detect and diagnose through monitoring power output and other SCADA systems, failures associated with combiner boxes and more granular hardware are much more subtle. Faults at this level of the plant, such as string outages and tracker hardware faults, when considered individually, contribute a less than one percent loss at the inverter level, and no more than a five percent loss at the combiner box level. While sensor suites are often installed on combiner box hardware, the current best practice to detect these faults is the performance of aerial IR scans, which are typically performed on an annual basis. While individual impacts from these faults are small, due to the low rate of detection these faults can accumulate across a PV plant and in total lead to more than 3% loss in power production [5].

To resolve these types of performance loss issues, utilities operate monitoring and diagnostics (M&D) centers to identify when hardware performance is not meeting expectations and inform business decisions regarding how best to address these issues. Recent technological trends have led to the adoption of new tools that utility M&D centers can use to better detect more subtle performance loss issues through machine learning. Because many utilities employ several different modes of power generation, with vastly different physics, sensors, and supporting hardware, M&D software solutions that are broadly applicable to a wide array of industrial systems become attractive choices for adoption.

To support such a diverse selection of systems, these solutions often use purely data-driven approaches and lack any sense of physical intuition regarding the system that they are being used to model. To keep the tools user friendly, they often employ simple machine learning techniques, such as clustering algorithms, to model asset performance. At utility M&D centers, these tools are used to deploy thousands of models across their power generating fleet; as a result, ease of use is a must. For a 100+ MW PV plant, it is not unreasonable that hundreds of models could be deployed for that one site alone.

Recent work has demonstrated the ability of physics-based models to accurately detect string outages and tracker faults using combiner box current data [6], [7]. An important step in evaluating the effectiveness of such models is to compare their performance to standard industry tools that are used for similar tasks. In this work, fault detection capabilities of the physics-based approach are compared against similarly constructed models in a state-of-the-art, industry standard M&D tool.

II. Methodology

A. Physics-Based Model

Recent works have led to the development of new physics-based approaches to aid in anomaly detection for utility M&D centers. One model uses Sandia National Lab's pvlib-python to perform accurate calculations of estimated plant power output based on detailed knowledge of the plant layout and hardware and measured weather data from the site [8]. Such approaches enable robust data filtering strategies, including the calculation of and filtering on solar elevation angle, as well as the deployment of cloud detection algorithms. More detailed information on the application of this physics-based model can be found in [6].

To aid in the deployment of physics-based models for utility-scale PV plants, model construction has been automated using machine learning techniques. This technique involves using DuraMAT's PV-Pro tool to estimate site architecture directly on a signal-by-signal basis [9]. Additional processing is performed using machine learning to estimate the number of modules per string and the unique values of strings per combiner across the entire site, to standardize and map sensor names to each piece of hardware, and to construct the necessary configuration information for the physics-based models. Detailed information about this process can be found in [7].

B. Advanced Pattern Recognition Model

Advanced pattern recognition (APR) software has become an industry standard monitoring and diagnostic (M&D) tool for electric utilities. To benchmark the physics-based models described in [6] and [7], new APR models were built for a PV utility site using an industry standard tool which uses a clustering algorithm to build performance models of industrial hardware. Since APR tools are typically intended to be widely applicable to many industrial assets, they often lack the functionality to do high fidelity data filtering and quality checks like those developed for the physics-based approach. Regardless, efforts were made to ensure that the APR models were developed with similar capabilities as the physics-based models, to the extent that it was possible to do so.

To build the APR models, one model was trained for each inverter to match the fault detection strategy of the physics-based approach, where signals are compared on an inverter-by-inverter basis. Each inverter model was trained on the inverter voltage and all individual combiner box currents, for that given inverter. No data filtering was implemented within the APR software at this stage, aside from basic signal quality checks.

The APR models yield two types of outputs when a particular operating point is identified as anomalous: (1) a model residual (per inverter), and (2) a percentage contribution, per signal in the model. The time frequency of the outputs is tunable for each model, and to keep evaluation time low, an interval of five minutes per evaluation was chosen. Since the physics-based approach evaluates anomalies on a day-by-day basis, these APR outputs were aggregated to perform a similar evaluation. At each timestep, the model residual was multiplied by the contribution of each signal to recover a residual for each signal in the model at each point in time. To calculate a daily alarm metric, these per-signal residuals were then averaged into a daily value. Finally, a flag was raised when this daily alarm metric exceeded a certain threshold. A particular combiner box was not considered faulted unless at least eight flags had been observed over a ten-day window. The flag threshold was varied across the dataset, and true and false positive rates were evaluated at each threshold value to evaluate performance of the APR model. The APR models were trained on approximately eight months of data.

C. Validation

An aerial IR scan was used as a ground truth model for which faults were currently present at the site. Both the physics-based and APR models were evaluated for a 90-day period, during which an aerial scan was performed at the end of that period. Detection metrics were analyzed based on the fault detection results on the day of the aerial scan.

III. Results

To validate the fault detection approaches, both models were evaluated on historical datasets that overlapped with the occurrence of an aerial IR scan. These aerial scans are assumed to be "ground truth," and allow for the calculation of true positive rates (TPR), false positive rates (FPR), and F1 scores for each model.

Fig. 1 shows initial results for both models that were executed for a single site. The APR model results shown here reflect the utility's default model parameters.

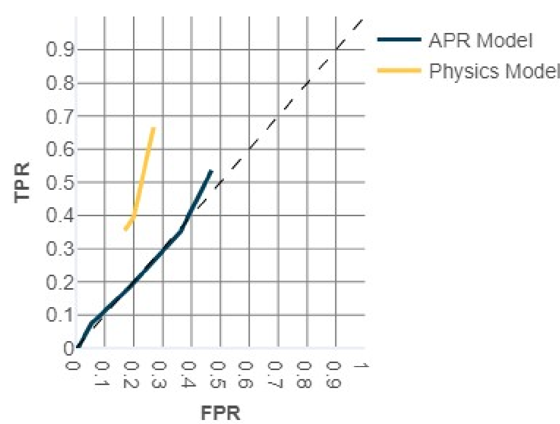

Fig. 1. Initial results comparing physics-based model anomaly detection to APR model anomaly detection at different detection sensitivities.

The initial APR model had effectively equal true and false positive rates, regardless of the detection threshold used. Using this model, if an alarm was raised for a particular hardware component, then it would be equally likely that that alarm was either a true or false positive. Essentially, the model is no more accurate than a coin flip in its ability to identify the presence of a fault, given that an alarm was raised. The partner utility identified that this model performed worse than they would have typically expected, and that such performance would warrant refinement of the model.

In this situation, the utility shared that they would usually refine the models by adjusting the number of clusters included in the model. Further discussion with the utility identified what additional strategies could be used in the models to mimic some of the data filtering approaches that are applied in the physics-based models. Sophisticated calculations, such as a cloud detection filter, were ruled out, but the utility identified that a filter on specific sensor values would be simple to implement within the software. To aid in improving the APR models, a filter was deployed to remove data where POA measurements were less than 200 W/m^2, effectively removing data points with particularly large cloud impacts. With these updates, the models were retrained and executed on the same data to recalculate the detection metrics. These results are shown in Fig. 2.

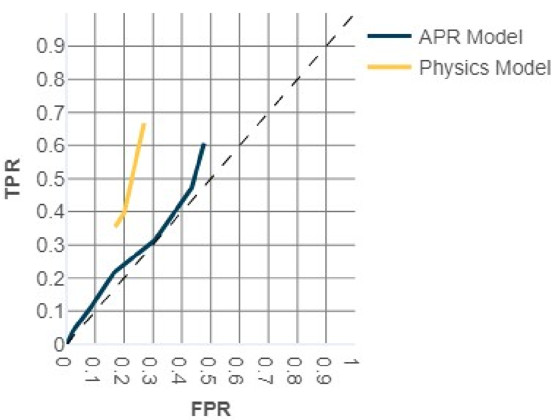

Fig. 2. Comparison of fault detection results from the updated APR models to the physics-based approach.

While the model modifications did have an impact on the detection metrics, the improvement in the detection accuracy was very slight. One reason for the higher detection rates in the physics-based approach is that the physics-based model does not rely on any training data while the APR model does. The physics-based model is able to make physics-driven calculations at each point in time based on site architecture and measured weather conditions, reducing reliance on any past operation. Conversely, the APR model is only as good as the data it was trained on. When there is normal operation that was not within the APR models training set, that data will look anomalous to the model even though the behavior is as expected. The physics-based model also allows for the calculation of derivative performance metrics that can be used to featurize the system and better draw out performance differences relative to normal operation. On the other hand, the APR system is limited to transformations of the recorded data. The APR model cannot make any inferences about the physics at hand beyond what can be calculated directly from measured data. Additionally, this model does not have any inherent knowledge of the site architecture.

Future work is planned to build similar models for a second site, which has a significantly larger hardware footprint. This second site has multiple aerial scans available for validation and will provide further insight into the difference in model performance across different site layouts.

Acknowledgements

This material is based upon work supported by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Solar Energy Technologies Office under Award Number DE-EE-0008976.

The authors thank William Hobbs, Southern Company, for assisting with this work.

References

[1] D. Feldman, K. Dummit, J. Zuboy, B. Smith, D. Stright, M. Heine, R. Margolis, "Fall 2023 Solar Industry Update," NREL, 2023.

[2] T. Hess, K. Tsai, "We expect solar will supply almost all growth in U.S. electricity generation through 2025," *Today in Energy*, U.S. Energy Information Administration, 2024.

[3] J. DeVilbiss, K. Antonio, "EIA expects U.S. annual solar electricity generation to surpass hydropower in 2024," *Today in Energy*, U.S. Energy Information Administration, 2024.

[4] "Solar Generation Index," kWh Analytics, 2022.

[5] "Global Solar Report: 2023 Edition," Raptor Maps. 2023.

[6] S. Sheppard, T. Cook, D. Fregosi, C. Perullo and M. Bolen, "Field Experience Detecting PV Underperformance in Real Time Using Existing Instrumentation," 2022 IEEE 49th Photovoltaics Specialists Conference (PVSC), Philadelphia, PA, USA, 2022, pp. 0307-0313, doi: 10.1109/PVSC48317.2022.9938522.

[7] S. Koskey, S. Sheppard, C. Teasley, C. Perullo, J. Kee, D. Fregosi, W. Li, "Photovoltaic Site Architecture Estimation Using Performance Data," 2023 IEEE 50th Photovoltaic Specialists Conference (PVSC), San Juan, PR, USA, 2023, pp. 1-6, doi: 10.1109/PVSC48320.2023.10359804.

[8] William F. Holmgren, Clifford W. Hansen, and Mark A. Mikofski. "pvlib python: a python package for modeling solar energy systems." Journal of Open Source Software, 3(29), 884, (2018). https://doi.org/10.21105/joss.00884.

[9] DuraMAT, Berkeley, CA. 2022. *PV Production Tools (PV-Pro)*, ver 0.0.4.