Field Experience Detecting PV Underperformance in Real Time Using Existing Instrumentation

Scott Sheppard[1], Tim Cook[1], Daniel Fregosi[2], Christopher Perullo[1], Michael Bolen[2]

[1]Turbine Logic, Atlanta, Georgia, 30308, USA
[2]Electric Power Research Institute, Charlotte, North Carolina, 28262, USA

***Abstract*—Maintenance at large-scale photovoltaic plants employs a mix of preventative and corrective maintenance practices. Large outages, such as an inverter tripping offline, are often easy to detect. More subtle sub-inverter faults and failures can accumulate and go unnoticed for months or years. A software-based fault detection method has been developed to analyze commonly measured data from large-scale PV plants for more timely detection of subtle underperformance. The method has been demonstrated on eight datasets from large-scale plants with high accuracy of detection. Results are validated using aerial infrared scanning. String outages are detected with a true positive rate of 73 percent and tracker issues are detected with a true positive rate of 88 percent. The developed method can be uniformly applied to photovoltaic plants across a range of scales and configurations to assess performance, quickly detect underperformance, and determine the source and location of failures. The results inform and improve operations and maintenance at PV plants, ultimately aiding in improved affordability, reliability, availability, and resiliency of solar electricity.**

Keywords—Solar, Photovoltaics (PV), Data Analysis, Operations, Maintenance, Monitoring, Detection, Failures, Faults, Underperformance

I. Introduction

Monitoring and Diagnostics (M&D) is of increasing importance for all forms of power generating assets with respect to improved operations, maintenance, reliability, and affordability. Advances in data analytics and algorithms, affordable data storage and computational power, ease of information transfer and communication, and sensor and instrumentation capabilities and affordability all offer new capabilities and insights to plant owners and operators [1], [2]. Utilizing and integrating each of these technological advancements to their fullest potential is an active research area.

Maintenance at many photovoltaic (PV) plants utilizes a mix of preventative and corrective maintenance practices. Corrective (a.k.a., reactive, break-fix) maintenance resolves failures after they occur. Preventative maintenance involves performing tasks before equipment fails and/or becomes unsafe. When preventative maintenance occurs, it can be based on multiple factors, generally categorized as time-based (e.g., annual) or condition-based (e.g., predictive analytics). Developing the basis for condition-based maintenance often requires analyzing data from the plant. Detecting when major pieces of equipment fail, such as inverters, is straightforward based on data collected from existing, intrinsic sensors embedded by the equipment manufacturer [3].

Detecting more subtle failures that occur within the DC collector field is more difficult. Failures that impact a single module, or even an entire string, are often undetectable through monitoring production yield alone. These failures can still be identified through various field inspection techniques which are not conducive to continuous online monitoring [4]. Some manual electrical testing methods, such as field I-V curve tracing, may take days or weeks to complete. Varying meteorological conditions over this time period can often mask any issues that may be present [5].

Time-based aerial infrared (IR) imaging is one common detection method used to detect failures in the collector field. This is usually performed on an annual cadence. Failures detected include electrical and mechanical issues such as module underperformance, module hot spots, broken modules, string outages, and broken or misaligned trackers (if applicable). While it depends on the configuration of the specific plant, these types of failures each individually represent a loss of less than 1% of the power output of a single inverter. The ability to supplement physical inspection with continuous monitoring based on available data would greatly improve plant reliability and affordability. The goal of the present work is to leverage existing sensor suites and data at large-scale PV plants to aid in operation and maintenance by improving the ability to detect and diagnose subtle failures in relative real-time. The focus is on the DC collector side of the plant where there is an abundance of data available for analysis and the faults and failures are typically more subtle and harder to detect than those impacting central inverters.

II. Monitoring and Diagnostic Challenges

Due to the variability in site architecture at PV plants, developing a standardized M&D approach can be very difficult. The numbers of modules, strings, combiner boxes, and inverters vary; panel and inverter hardware vary; solar irradiation and other weather varies; racking (e.g., fixed, single-axis tracking) and orientation varies. Determining how a plant should perform from one day to the next becomes an onerous M&D activity that is often outside the expertise or time constraints of most utility M&D centers. Considering current, voltage, and tracker channels (if applicable) at a large-scale PV plant, monitoring centers are typically evaluating signals from tens of thousands of sensors or more.

Many subtle issues go unnoticed either indefinitely or until they become severe enough that they must be fixed. Due to the physical size and remote location of many PV plants, routine inspections are not performed frequently as they require a large amount of employee hours to perform [5]. As a result, even when degradation is expected, locating the source of the power loss to a specific section of the plant can be a nontrivial task. Periodic IR and visual overscans help in this effort, but as they are often performed on an annual basis many failures go undetected for weeks to months at a time.

The developed fault detection method addresses these challenges by comparing measured data from the plant to performance calculated by a physical model of the plant coupled with feature extraction. This modeling approach provides the opportunity to leverage all the data currently being measured to find periods and locations of underperformance. While the proposed method is still reactive, it moves maintenance towards an analytics-driven, condition-based approach.

III. Fault Detection with PV Plant Data

A. Performance Model Framework

Fig. 1 shows the general workflow of the developed fault detection model. There are three key components to the model: the physical plant model, cleanup of measured data, and comparative analysis between the model and measured performance and among similar components at the plant. The physical models have been developed using pvlib-python, an open-sourced software administered by Sandia National Lab [6], [7]. This software package was chosen due to its extensive and robust hardware library that allows for accurate replication of a commercially operating plant. Additionally, the pvlib-python package allows the user to specify plant configuration information, such as the number of modules per string and number of strings per array, which allows for performance calculations specific to the user defined system. When site specific weather data is available, as is the case at most large-scale PV plants, the models can be used to calculate expected current, voltage, and power for each modeled subarray. This functionality allows for site specific modeled performance that accounts for weather conditions experienced at the plant rather than relying on clear sky models to calculate expected performance.

B. Data Cleanup

The next key step of the analysis process is cleanup of the measured data. The following simple rules were arrived at through iteration with the models, are easy to apply, and reduce false alarms:

- Remove points with plane of array irradiance below 500W/m^2
- Remove points with solar elevation below 30°
- Remove known bad data points

These filters have the greatest impact when applied in tandem with each other. The first two filters primarily remove points early and late in the day (and overnight). At low solar elevation, adjacent rows tend to self-shade which adds a large degree of noise to the data for these time periods. Removing known bad values from the dataset also significantly reduces noise in the data. These values are typically values that are outside the bounds of the sensors measuring capabilities and any null readings that come from any of the sensing hardware (for example, many data historians have a “time out” value that indicates that no data is being logged). Other data quality filters than can be applied are detailed in [8].

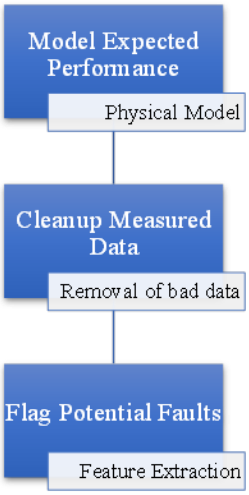

Fig. 1. Physical model framework for detecting faults at PV plants.

Typically, inverter max power point trackers will shift the DC operating point of an array to match the power rating of the inverter hardware, resulting in DC clipping. As a result, the appearance of a fault is present in the DC power output of an array since the DC power output will be low relative to its expected output. If analysis is being done on the DC power output of an array, it is recommended that any points where the DC power is 99% or greater than the inverter rating be removed from analysis.

C. Cloud Detection

Another data cleanup step that is used is a cloud detection filter on the data. Transient cloud cover adds a significant amount of noise to PV data, and removing this data provides a much cleaner dataset for further analysis. This functionality is provided within pvlib-python and is described in [9] and [10]. The application of the pvlib-python cloud detection thresholds was modified to increase performance across differing site architectures.

Through further usage of the modified cloud detection algorithm from pvlib-python, the method proved to be too aggressive in its filtering, often filtering seemingly clear sky irradiance data as “cloudy”. The main driver for this overclassification of cloudy data was determined to be caused by a significant difference in the clear sky model of irradiance used by pvlib-python to perform cloud detection and the on-site measurements. An example of this difference is shown in Fig. 2. Due to this difference, a new method for filtering data based on cloud cover through comparison to historical, locally measured irradiance data was developed.

The new filter uses statistical methods and predictions based on a polynomial regression of historical plane of array (POA) irradiance measurements to filter out cloudy data. The

This work is funded in part by the U.S. Department of Energy Solar Energy Technologies Office, under award number DE-EE-0008976.

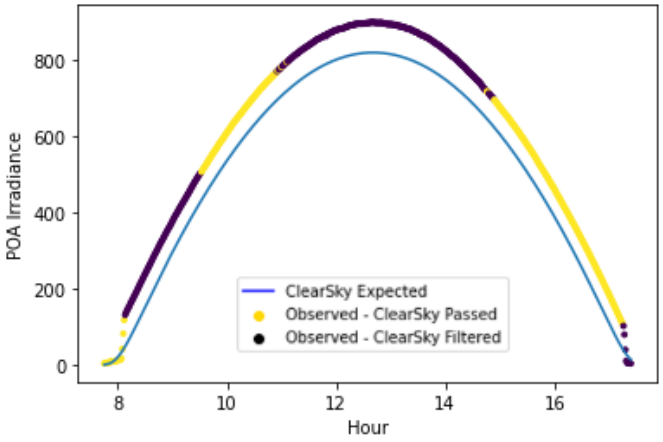

Fig. 2. Discrepancy between measured irradiance and the clear sky model.

architecture for the new clear sky filter is shown in Fig. 3. For a given day, the POA timeseries is first statistically filtered to eliminate points with erratic or rapidly changing POA signal. The underlying assumption is that POA measurements should vary smoothly with time in the absence of cloud cover. The remaining data is then regressed with a 4th-order polynomial to calculate a prediction of the clear sky irradiance. Depending on the amount of data removed from the first round of filtering, data from previous days is overlaid on the current day and included in the regression. Then, a second filtering step eliminates data points which deviate from both the predicted clear sky irradiance value and slope by a significant margin at each point in time, adding datapoints to the initial filter. Optionally, this prediction-filter step can run in a loop, adding points to the filter with each iteration.

To create the cloud detection filter for a given day, first, the derivative of the POA irradiance values is calculated for each datapoint over the course of a day. Then, the standard deviation of the derivative is calculated for a five-minute moving window centered on each time step. This standard deviation is normalized by the maximum POA value in the entire timeseries, and all points with a standard deviation greater than 1% of the maximum POA irradiance value are filtered out. This initial statistics-based filter is applied to each day in the dataset. The effect of this statistics-based filter on an example day is shown in Fig. 4.

For each day, this filtered irradiance data is used to create a prediction of clear sky POA irradiance based on a polynomial regression of data. The 4th-order polynomial is suitable because it can closely fit nearly any daily solar profile, regardless of location, tracking system, or orientation. The previous six days of filtered POA data is also included in the regression,

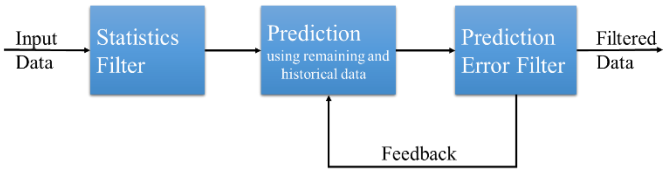

Fig. 3. Process for filtering measured POA data to create a data-based expected clear sky irradiance prediction.

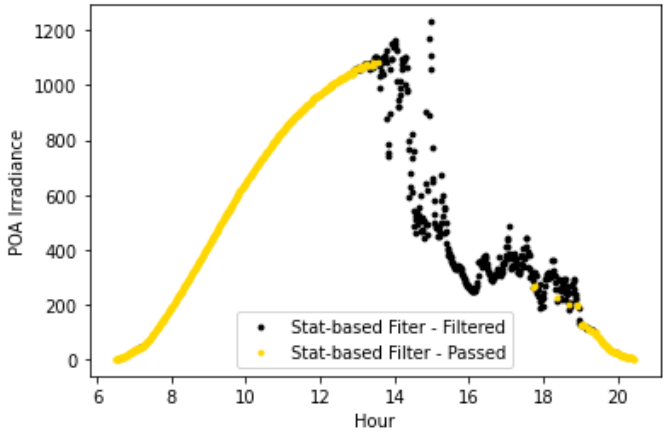

Fig. 4. POA irradiance for an example day highlighting the impacts of the initial statistics-based filter.

dramatically improving the accuracy and reliability of the prediction on cloudy days. Due to changes in sunrise and sunset from day to day, six historical days of data was chosen as the upper limit for the time window for data inclusion to keep the endpoints in the measured daily POA curves from drifting too far apart from each other. An exception is made for individual days in which more than 75% of data remains after the initial statistics-based filter. On these days, the regression uses only data from the current day to regress the 4th-order polynomial.

An example of the improvements to the regression provided by increasing the number of days in the regression is shown in Fig. 5. This figure shows three regressions overlaid on the filtered data shown in Fig. 4. Each line represents a regression model including differing amounts of historical data in the irradiance regression. As Fig. 5 shows, a larger amount of data leads to a more accurate prediction of clear sky irradiance.

The trained 4th-order regression model is used to create a prediction of the clear sky POA for the current day. To further improve the predicted clear sky value, the filtered data is subject to an additional round of filtering, incorporating the results from the clear sky regression. First, the slopes of expected clear sky irradiance and measured irradiance are calculated over a nine-

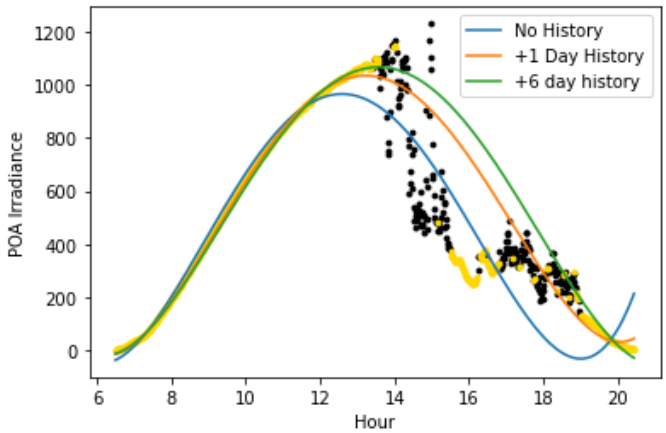

Fig. 5. Results of the 4th-order regression of POA irradiance as the number of days included in the regression increases.

minute interval, centered at each time step. The final filter assesses the difference in both values and slopes of the measured POA and the predicted clear sky POA at each time step. Values which fall more than 20% below the predicted value (normalized to the maximum value in the dataset) are removed. Time steps with a slope deviating by more than 50° from the predicted slope are also removed.

The remaining values in the timeseries are again regressed in a 4th-order regression model to update the prediction and iterate the second stage of filtering. Through testing of this clear sky modeling approach, it was found that only one iteration of the regression and 2nd stage of filtering was necessary for the clear sky filter to reach a stable state.

Fig. 6 shows the results of the prediction-based filter, with the red line approximating the 20% difference threshold for the difference in expected and measured POA values. The points in red represent those that passed the original statistics-based filter (shown in Fig. 4) but fall outside the thresholds for either predicted value or slope differences and were removed by the prediction filter. Near the peak of the curve, a small number of datapoints are eliminated based on the deviation in slope though they are still close to the predicted value.

This new filtering approach shows an improvement in performance over original detection method from pvlib-python. Comparison on a full year of POA data from a single site, counting only daytime data, shows that the two filters yield the same result (pass or remove) for 77.5% of the entire series. The filter results differ on the remaining data points, with 16.2% of data points filtered only by the pvlib-python approach, and only 6.3% of data points remove by the new filtering approach. These results are listed in Table I. A significant portion of the data filtered only by the pvlib-python method represent inappropriate filtering of clear sky irradiance values. This is observable in Fig. 7a for a relatively clear day, in which the green datapoints represent datapoints which were removed by pvlib-python filter but passed the new filter. Similarly, many of datapoints that were not filtered by the pvlib-python method should have been removed. This is most often the case for short duration cloud events, exemplified as the blue points in Fig. 7b.

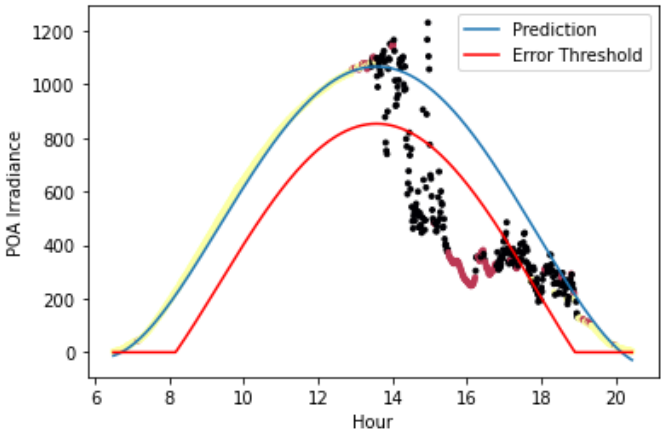

Fig. 6. The results of the additional filtering on expected cloudy points.

TABLE I. SUMMARY AND COMPARISON OF CLEAR SKY FILTERS APPLIED TO DAYTIME DATA.

		New Filter Classification		
		Cloudy	***Clear***	
Pvlib-python Filter Classification	***Cloudy***	28%	16.2%	**44.2%**
	Clear	6.3%	49.5%	**55.8%**
		34.3%	**65.7%**	

The 4th order polynomial regression provides a significantly better expected clear sky POA signal based on historical data. This result greatly reduces false positive filtering of clear datapoints and provides more available data, post-filtering. The regression-based filter also slightly improves filtering of short duration cloud events.

D. Feature Extraction

Finally, feature extraction analysis is performed between cleaned, measured data and the expected, model prediction. The pvlib-python model provides estimated values of expected operating voltage, current, and power as well estimates of each subsystem's short-circuit current (Isc) and open-circuit voltage (Voc). This information can be used to construct a rough

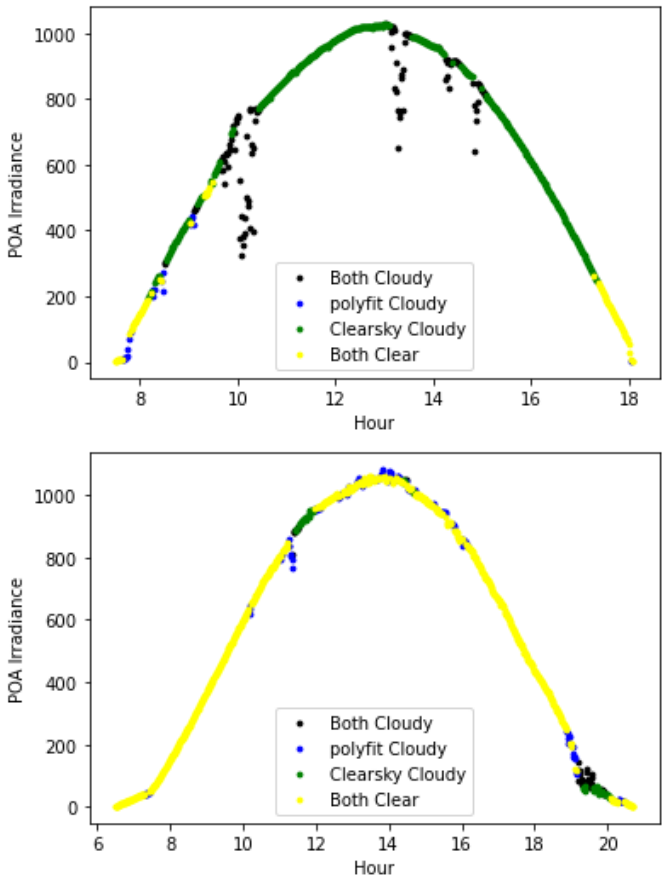

Fig. 7. Sample results comparing the original pvlib-python cloud filter to the new filter for a) a day where the original filter misclassified a majority of the data as cloudy and b) a day where the original filter missed some periods of cloud cover.

approximation of the expected I-V curve of the modeled subsystem. As different hardware failures have real impacts to the subarray's measured current or voltage (or both), it follows that the shape of the subarray's I-V curve would be impacted as a result of the failure [11]. The two sources of information are used to calculate derived metrics that are used in further anomaly detection. These derived metrics are compared among similar hardware components, where any detected anomalies are assumed to be indicative of a fault present at the PV plant. This analysis is performed on a day-by-day basis to ensure that adequate data is available to perform the necessary fault detection. This fault detection method has been applied to several datasets from many plants of varying sizes and geographic locations with success.

IV. Results

A. Historical Datasets

Historical data was provided for six different large-scall PV plants across the southern US, representing a wide range of plant architectures. These plants covered a broad range of several defining characteristics, including:

- Location – southeastern US to southwestern US
- Power generating capacities – 10+ MW to 200+ MW
- Number of inverters on site – 20+ to 120+
- Number of combiner boxes per inverter – <10 per inverter to 20+ per inverter
- Number of strings per combiner box – <10 to 90+ per combiner box
- Specific inverter hardware
- Specific module hardware

Each of these sites has single-axis trackers across the DC collector field.

The data provided for each site consisted of an extremely large number of data channels, spanning several months to more than a year of data per site. The following measurements were extracted from site historians and were available for fault detection analysis:

- Inverter DC power, current, and voltage for all inverters on site
- Combiner box DC current for all combiner boxes on site
- Meteorological measurements (POA irradiance, ambient temperature, and wind speed) for all meteorological stations at each site
- Tracker positions for all tracker controllers at each site

For each plant, a physical model was built to match the site architecture to calculate expected performance using the meteorological data provided at each plant. The modeled performance data was then compared to measured data at each combiner box to perform the fault detection analysis.

B. Validation

To validate the fault detection method, historical data was extracted from each plant for time periods concurrent with previously performed aerial IR and visual scans. These aerial scans were performed by a third party to notify the plant of failures impacting their hardware. These scans identified a wide range of failures, including tracker-related failures, sectional, string, and partial string outages, and a number of failures that singularly impact individual modules. These failures were aggregated to the combiner box level to match the granularity of sensor data that was available at each site. Each combiner box was then labeled by the failure that was anticipated to have the most impact to the sensor measurements. In general, tracker-related failures have the greatest impact to power input, followed by sectional and string outages, while module and sub-module failures will have the smallest impact to the combiner box measurements.

For validation purposes, it was assumed that the failures found by the aerial scans were completely correct, both in terms of failures being present and absent. Each dataset was processed through its respective model to calculate expected current and voltage measurements for each combiner box at the plant. The potential faults identified by the physical model were then validated against the failures identified by the aerial IR scans to determine true and false positive rates. Two of the sites had scans available for two consecutive years. In all, this allowed for eight different validation datasets.

C. Confusion Matrix Development

For each plant, the respective fault detection model was run on site data concurrent to the aerial scan to determine if a fault may be detectable at the combiner box level using historian data alone. The results from the fault detection model were then joined to the results from the aerial scans. Using the fault detection results and the aerial scans, a confusion matrix was constructed for each dataset to assess the accuracy of the fault detection model. For the purposes of this paper, the confusion matrix results are defined as follows:

- True positive – both the aerial scans and the fault detection model noted a fault, within a given combiner box
- False positive – the fault detection model noted a fault while the aerial scans did not, within a given combiner box
- True negative – neither the aerial scans nor the fault detection model noted a fault, within a given combiner box
- False negative – the aerial scans noted a fault while the fault detection model did not, within a given combiner box

The fault detection model detects the presence of faults on a day-to-day basis. Sample results for a single inverter are shown in Fig. 8. This figure shows how the confusion matrix classification changes for each combiner box over the course of a week. One observation is the generally steady state of the fault detection results; the model does not repeatedly bounce between faulted and unfaulted states. This time component has not been

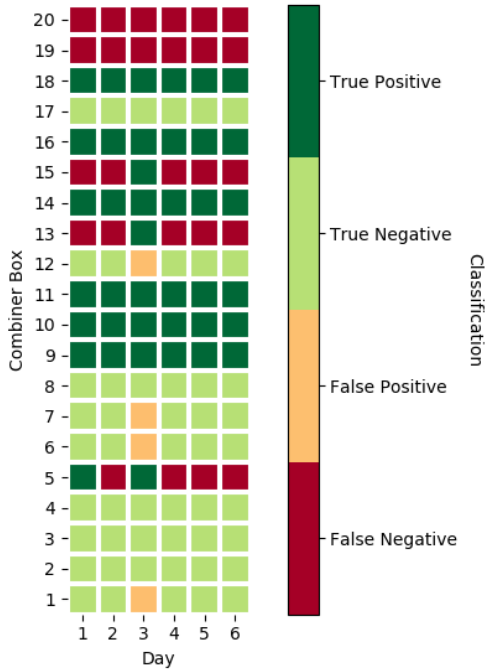

Fig. 8. True positive, true negative, false positive, and false negative classifications for fault detection for a single inverter. The aerial scan was performed on Day 5.

heavily analyzed to-date but could be used as part of the detection model to determine the presence (or lack) of faults. Additionally, since the presence of faults are evaluated on a day-by-day basis, this analysis method allows for the detection of the onset of faults. If implemented in a near real-time environment, this method could notify operators of faults as they happen, an improvement over the typically annual aerial inspections.

Since the aerial scans only reflect the state of the hardware on a single day, the remaining analysis in this paper is focused only on the fault detection results on the day of the scan. A potential next step is to consider how the time-based change in fault detection could be included to improve the accuracy of the fault detection model.

D. Fleet Summary

For each dataset, the results of the fault detection model were joined with the results of the aerial scans. To characterize the types of failures that were accurately detected by the model, the aerial scan results were categorized into three types: minor failures, string outages, and tracker failures. A listing of the failures in each of these categories is shown in Table II. In short, minor failures impact one to a few modules, string outages impact one to a few strings, and tracker failures are failure of individual tracker controllers and impact multiple strings.

The results from all datasets are summarized in Table III. Note that several sites in the summary are lacking tracker failures. Depending on the site, this is due to either a lack of labeling of tracker failures or a lack of tracker failures.

As shown in Table III, this fault detection method has consistently shown true positive rates at 75% or higher for string outages when validated against aerial IR scans performed the

TABLE II. SUMMARY OF LABELS APPLIED TO EACH FAILURE.

Summary Category	***Failures***
Minor	Cell Faults
	Hot Spots
	Module Anomalies
	Sub-module Anomalies
String Outage	String Outage
	Sectional Outage
Tracker	Tracker at stow
	Tracker anomaly

TABLE III. FAULT DETECTION TRUE AND FALSE POSITIVE RATES.

	True and False Positive Rate Summary		
Site	***Fault Type***	***TPR***	***FPR***
1	Normal	-	36%
	Minor	34%	-
	String	100%	-
	Tracker	77%	-
2a	Normal	-	39%
	Minor	44%	-
	String	88%	-
	Tracker	74%	-
2b	Normal	-	44%
	Minor	44%	-
	String	90%	-
3	Normal	-	67%
	Minor	86%	-
	String	67%	-
4	Normal	-	32%
	Minor	41%	-
	String	54%	-
	Tracker	56%	-
5a	Normal	-	7%
	Minor	92%	-
	String	88%	-
5b	Normal	-	30%
	Minor	33%	-
	String	91%	-
6	Normal	-	50%
	Minor	82%	-
	String	92%	-

same day. Depending on the site architecture, these faults individually represent approximately a 1% loss of power output at the inverter level and are the types of faults that are small enough to go unnoticed when monitoring inverter and plant power outputs and, when considered in aggregate, are one of the largest sources of power loss detected by aerial overscans. The fault detection method shows good detection rates for these faults, but, unlike aerial scans, is able to perform fault detection analysis on an ongoing basis. An overall summary of the fault detection method is shown in Table IV. The overall true positive rate across all faults and all eight datasets was 66%. The true positive rates for the faults with the largest impact to plant power production are significantly higher: 88% for string outages and 73% for tracker faults. Even the minor faults, which only impact a single module, where detected with approximately a 50% true positive rate. With additional tuning of internal detection thresholds, these results are expected to improve.

V. Conclusions

The results discussed above show the strength of using a physical model in conjunction with real PV data to perform fault detection on DC collector field data. Currently, many sites only perform aerial scans annually, which limits their insight into the presence of subtle, small-scale faults. The developed fault detection method operates on data that is typically available at M&D centers on a continuous basis and is capable of providing reliable fault detection year-round. Additionally, the detection locates faults to specific combiner boxes, reducing the field inspection effort of locating the fault when maintenance needs to be performed. With real time results, M&D centers could leverage this fault detection model to schedule maintenance as issues arise, rather than wait for the next annual scan to notify them of hardware failure. Production, and therefore revenue, is reduced when these failures are undetected and uncorrected.

To ensure that operators do not become burdened by attempting to diagnose false alarms, future work is aimed at reducing the false positive rate while keeping the true positive rate high. One approach will be to consider how day-to-day variations in the fault detection results can be leveraged to determine if a fault is truly present. Additionally, more tuning of the internal fault thresholds could lead to potential improvements in true and false positive rates.

Table IV. Overall Summary of the Fault Detection Model.

Fault Type	***TPR***	***FPR***
Overall	66%	28%
Normal	-	28%
Minor	45%	-
String	88%	-
Tracker	73%	-

References

[1] "Chapter 2: Energy sectors and systems," Quadrennial Technology Review 2015, US Department of Energy, 2015.

[2] "Chapter 4: Advancing Clean Electric Power Technologies," Quadrennial Technology Review 2015, US Department of Energy, 2015.

[3] W. Brooks, T. Basso, and M. Coddington, "Field guide for testing existing photovoltaic systems for ground faults and installing equipment to mitigate fire hazards," United States, https://doi.org/10.2172/1225963.

[4] R. R. Hill, G. T. Klise. And J. R. Balfour, "Precursor report of data needs and recommended practices for PV plant availability, operations, and maintenance reporting". https://osti.gov/search/identifier:1169447.

[5] Walker, H. A. "Best practices for operation and maintenance of photovoltaic and energy storage systems; 3rd Edition," United States. https://doi.org/10.2172/1489002.

[6] W. F. Holmgren, C. W. Hansen, and M. 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.

[7] W. F. Holmgren et al, pvlib/pvlib-python: v0.7.2 (v0.7.2). Zenodo. https://doi.org/10.5281/zenodo3762635.

[8] "Photovoltaic systems performance – Part 3: Energy evaluation method," IEC Technical Specification, IEC TS 61724-3, ISBN 978-2-8322-3531-7.

[9] M. Reno and C. Hansen, "Identification of periods of clear sky irradiance in time series of GHI measurements," Renewable Energy, vol. 90, 2016, pp. 520-531.

[10] B. Ellis, M. Deceglie and A. Jain, "Automatic detection of clear-sky periods using ground and satellite based solar resource Data," 2018 IEEE 7th World Conference on Photovoltaic Energy Conversion (WCPEC) (A Joint Conference of 45th IEEE PVSC, 28th PVSEC & 34th EU PVSEC), 2018, pp. 2293-2298.

[11] M. Kontges, et al. "Performance and reliability of photovoltaic systems, subtask 3.2: Review on failures of photovoltaic modules," IEA Photovoltaic Power Systems Programme, Task 13, 2013, ISBN 978-3-906042-16-9.