Automatic Detection and Analysis of Photovoltaic Modules in Aerial Infrared Imagery

Sergiu Dotenco[1], Manuel Dalsass[2], Ludwig Winkler[3], Tobias Würzner[3], Christoph Brabec[2], Andreas Maier[3], and Florian Gallwitz[1]

[1] Department of Computer Science,
Nuremberg Institute of Technology, Nuremberg, Germany
[2] ZAE Bayern, Germany
[3] Pattern Recognition Lab, Friedrich-Alexander University Erlangen-Nürnberg,
Erlangen, Germany

Abstract

Drone-based aerial thermography has become a convenient quality assessment tool for the precise localization of defective modules and cells in large photovoltaic power plants. However, manual evaluation of aerial infrared recordings can be extremely time-consuming. Therefore, we propose an approach for automatic detection and analysis of photovoltaic modules in aerial infrared images. Significant temperature abnormalities such as hot spots and hot areas can be identified using our processing pipeline. To identify such defects, we first detect the individual modules in infrared images, and then use statistical tests to detect the defective modules. A quantitative evaluation of the detection and analysis pipeline on real-world infrared recordings shows the applicability of our approach.

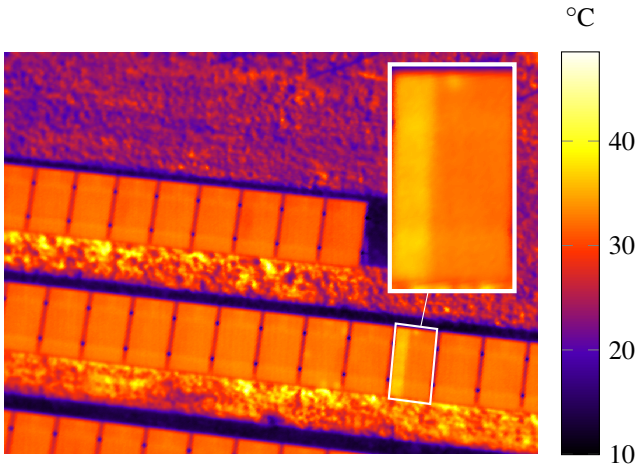

Figure 1. Radiometric aerial infrared image data of a photovoltaic power plant displayed in false color. The left substring of the marked photovoltaic module is inactive.

1. Introduction

Over the last decade, the market for photovoltaic systems has grown tremendously. In Germany alone, about 1.5 million photovoltaic systems with an overall peak power of approximately 38.5 GW were installed at the end of 2014 [1]. Considering the necessary continuous maintenance of such a high number of installed photovoltaic modules, there is a need for efficient quality assessment tools.

Particularly in recent years, operators and investors of photovoltaic power plants started appreciating the benefits of infrared thermography. Infrared imaging is time-saving, non-destructive, and contactless. These features allow an easy data acquisition under real operating conditions. In infrared images of photovoltaic modules, temperature abnormalities such as hot spots and hot areas indicate irregularities, which are often caused by defects in photovoltaic modules [2, 3, 4, 5]. In combination with remote-controlled unmanned air vehicles (UAVs), infrared thermography is therefore an especially convenient quality assessment tool for the precise localization of defective photovoltaic modules and cells in large photovoltaic power plants [6, 7]. Fig. 1 shows radiometric aerial infrared image data of a solar power plant acquired using a drone.

Despite the technical progress, aerial infrared imagery is still analyzed mostly manually, which can be extremely time-consuming. With regard to future monitoring and early detection systems for photovoltaic power plants, this work aims at bringing forward the automatic evaluation of aerial infrared recordings.

1.1. Contributions

In this work, we propose an image processing pipeline for the automatic evaluation of aerial infrared low-resolution images of photovoltaic power plants. The pipeline consists of two key components.

1. Detection of individual photovoltaic modules within an infrared image.
2. Analysis and detection of malfunctioning photovoltaic modules using infrared thermography.

Detected defects include, for example, hot spots and hot areas within the glass surface area of an identified photovoltaic module.

1.2. Outline

The remainder of this paper is organized as follows. In Section 2 we discuss related work. Section 3 provides details on our experimental setup and data acquisition. Sections 4 to 5 describe our photovoltaic module detection and analysis pipeline. In Section 6 we present our experimental results. A conclusion and an outlook on future work is given in Section 7.

2. Related Work

Infrared thermography is a commonly used technique for various tasks. Dios *et al*. [8] for instance, equipped a UAV with infrared cameras to automatically detect heat loss in buildings through windows and uninsulated walls. In their approach, areas with thermal values higher than those of their surroundings are segmented. They define a specific temperature threshold, which is, however, highly dependent on the surrounding air temperature. The measured temperature distributions are modeled both using a single Gaussian distribution in case of uniform emissivity, and two Gaussian distributions in the non-uniform emissivity case.

Tsanakas *et al*. [9] used Canny edge detection to find hot spot formations in photovoltaic modules by means of infrared thermography. The objective of their work was to show that infrared imaging can be used for detection of malfunctioning photovoltaic modules. Their approach is able to identify hot spot formations that were diagnosed to specific defective cells. However, the detection algorithms lacks a defect classification ability and its performance highly depends on the environmental conditions, which must be taken into account during field measurements.

Leotta *et al*. [10] developed an image processing tool, which can automatically detect single thin-film photovoltaic modules in infrared images. Using a hand-held infrared camera focusing on one photovoltaic module in the field, they can achieve a detection error of less than 5 %. However, their software could only identify up to 30 % of the photovoltaic modules of their aerial infrared imagery which had a spatial resolution of 382 × 288 pixels and was taken from a flying altitude of 20 m. Thus, their image processing tool is not suitable for the automatic inspection of photovoltaic modules in aerial infrared images yet.

Figure 2. Remote-controlled octocopter with an integrated infrared and RGB camera, which we use for data acquisition.

3. Drone-based Infrared Thermography

The radiometric data was generated using drone-based infrared thermography. The basis of our measurement system is a remote-controlled octocopter with an extensive navigation system, the DaVinci Copters ScaraBot X8 [11] shown in Fig. 2. Two lightweight cameras are installed beneath the octocopter, the infrared camera Optris PI450 and the RGB camera GoPro Hero3+.

Test objects were two photovoltaic power plants in northern Bavaria, Germany, with an overall peak power of approximately 7 MW_p. The test objects are installed on flat terrain and consist of polycrystalline photovoltaic modules having 60 solar cells and three bypass diodes (each for 20 solar cells on the left side, the center, and the right side, respectively). The photovoltaic modules are installed at a tilt angle of 25°.

The flight path of the octocopter with a flying altitude of approximately 20 m above the test objects had a systematic meandering pattern and was carried out by means of way point flight software. With regard to easy-to-handle infrared recordings for a future fully automatic analysis algorithm, the observation angle of the cameras was as orthogonal as possible to the flat ground of the photovoltaic power plant. However, the observation angle of the infrared camera towards the glass surface of the photovoltaic modules was less than 30° as recommended in [12].

The measurements were conducted in March and April 2015 on calm, sunny days around noon and in the afternoon. The irradiation values during this time ranged from 600 W m^{-2} to 1000 W m^{-2}. The sky was cloudless without any reflections of the sky on the photovoltaic modules in the infrared recordings. For evaluation purposes, we extracted infrared images at the original resolution of 382 × 288 pixels from recorded video sequences (see Fig. 1).

4. Detection of Photovoltaic Modules in Infrared Images

To analyze and assess the quality of photovoltaic modules with respect to defects or soiling, we first detect individual modules in infrared images. We follow a statistical, data-driven approach that consists of the following steps: 1) normalization, 2) thresholding, 3) orientation estimation of the photovoltaic modules, and a final 4) correction and refinement.

Typically, infrared images exhibit a low signal-to-noise ratio (SNR). However, compared to the background, photovoltaic modules irradiate considerably higher temperature, which also has a low local variability. By additionally assuming the temperature to be normally distributed, we can apply statistical methods to segment photovoltaic modules from the background. Although the assumption of photovoltaic modules being normally distributed is simplistic, our results clearly show that modeling the temperature using the Gaussian distribution allows to segment photovoltaic modules very accurately.

4.1. Normalization

As a first step, we normalize the temperature of infrared images to a constant range, since the temperature can heavily fluctuate. In our data, for instance, the temperature ranges from 0 °C to 57 °C. Normalizing the temperature allows us to define a set of parameters that works well for a variety of infrared images with varying temperature ranges.

For normalization purposes, we discard temperature values below a predefined threshold τ by setting the values to zero. By discarding low temperatures, we remove unnecessary information while preserving temperature close to the operating temperature. The actual normalization is performed using

$$T_{\text{norm}}(u,v) = \frac{\max\left(T(u,v) - \tau, 0\right)}{\max_{(u',v')\in\Omega} T(u',v') - \tau} \quad , \tag{1}$$

where $T\colon \Omega \to \mathbb{R}$ specifies the temperature in °C at the coordinate $(u,v) \in \Omega \subset \mathbb{N}^2$ within the infrared image T. $\max_{(u',v')\in\Omega} T(u',v')$ represents the largest temperature value in T, also given in °C. $T_{\text{norm}}\colon \Omega \to [0,1]$ defines the normalized temperature map. In our experiments, we use $\tau \coloneqq 20\,°\text{C}$, which removes shadowed grass areas. Fig. 3(b) shows a typical example of a normalized temperature map.

4.2. Automatic Thresholding

In the next step, we separate the modules from the background. We use automatic thresholding based on the idea by Rother *et al.* [13]. Their approach uses Gaussian mixture models (GMMs) for segmentation purposes. However, the approach does not perform segmentation completely automatically. It requires some user interaction for labeling foreground and background regions.

Instead of using GMMs, we model the foreground, *i.e.* the photovoltaic modules using a Gaussian distribution. We represent the probability density function (PDF) in terms of the Gaussian function

$$g(t \mid \alpha,\mu,\sigma) = \alpha \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right) \quad , \tag{2}$$

where t represents the temperature, α the amplitude, μ the mean, and σ the standard deviation. We denote the PDF of (2), *i.e.*, the normalized Gaussian function as $\tilde{g}$. The goal now is to estimate all three unknown parameters from the normalized infrared image T_{norm}.

4.2.1 Parameter Estimation

We approximate μ by the maximum temperature value $\bar{\mu}$ in the upper 80 % of the histogram with the bins $(b_t)_{t=0,\dots,n}$ computed from the normalized temperature map T_{norm} as

$$b_t \coloneqq \sum_{(u,v)\in\Omega} \begin{cases} 1 & \text{if } \lfloor c \cdot T_{\text{norm}}(u,v) \rfloor = t \\ 0 & \text{otherwise} \end{cases} \quad . \tag{3}$$

Here, c is a scale factor that maps normalized temperature values to histogram bin indices $t \in \{0,\dots,n\} \subset \mathbb{N}$.

Now, let $\bar{\mu}$ be the approximation of μ. By extracting 64 bins from the histogram in Eq. (3) around the bin $\bar{\mu}$ (32 bins to the left and right of $\bar{\mu}$, respectively), we estimate an initial probability mass function $\Pr(T = t) =: p_t$ of Eq. (2). Here, T refers to a random variable with the realizations t and the probability p_t. We then use nonlinear optimization to refine the estimated distribution parameters according to the objective function

$$(\alpha,\mu,\sigma) = \arg\min_{\alpha,\mu,\sigma} \frac{1}{2} \sum_t \left(p_t - \tilde{g}(t \mid \alpha,\mu,\sigma)\right)^2 \tag{4}$$

The initial value of α is set to the value of the bin corresponding to $\bar{\mu}$. The remaining parameters μ and σ are estimated from the histogram.

4.2.2 Segmentation of Photovoltaic Modules

To finally separate the foreground from the background, we use

$$\vartheta \coloneqq \mu - 4\sigma \tag{5}$$

as the threshold. This threshold works especially well for distributions with a small standard deviation σ. However, infrared images with a high negative skewness in the sampled distribution can cause parts of the rows to be omitted. We therefore address this issue in the final correction step.

4.3. Removing Background Clutter

Automatic thresholding does not completely exclude background regions from the foreground due to their high variance. Hence, to identify remaining background regions in the foreground mask, we estimate the variance T_{var} of the normalized temperature map T_{norm} using a 3×3 kernel. This allows to detect regions that have a high variability and are likely to be part of the background.

Subsequently, we perform histogram normalization on the variance map T_{var}, which spreads its values evenly. By applying a binary threshold at two thirds of the T_{eqvar} spectrum, we detect regions of high variance and eventually exclude them from the thresholded map T_{norm} while leaving homogeneous regions as part of the foreground mask $\mathcal{M}_{\text{rows}}$ consisting only of photovoltaic module rows.

Note that variance map estimation must be used in conjunction with automatic thresholding, since the variance map alone would not exclude shadows of photovoltaic modules, which have a low variance.

4.4. Photovoltaic Modules Row Orientation Estimation

$\mathcal{M}_{\text{rows}}$ contains both individual photovoltaic modules and groups of connected modules. The latter must be separated into individual modules. For this purpose, we estimate the orientation of photovoltaic module rows, and rotate the temperature map into an upright orientation.

We first create a binarized version of $\mathcal{M}_{\text{rows}}$, and apply morphological closing using a 3×3 kernel. To disambiguate between the rows, we use the Watershed transform [14]. $\mathcal{M}_{\text{rows}}$ is used as the marker. After labeling, we can operate on each row individually.

We proceed by extracting the contours of the resulting individual rows. Then, we determine the photovoltaic module row having the largest area, and fit an ellipse to it. The orientation of the row corresponds to the orientation of the fitted ellipse. Alternatively, one can use a histogram based voting scheme to determine the angle of most dominant lines as in [15].

4.5. Extraction of Individual Photovoltaic Modules

Once the rows in $\mathcal{M}_{\text{rows}}$ are brought into an upright position, we extract the individual modules. Here, we first apply a morphological opening with a 5×5 kernel to eliminate any remaining noise or outliers not belonging to the modules. Then, we use a morphological closing with a 1×7 kernel to close the gaps between module rows. Once again, we apply the Watershed transform to segment the rows, which are finally approximated by a rectangle fitted to them.

As individual modules are not connected to each other, regions that separate the modules can be identified by a significant drop in the temperature in each row. To detect these separating regions, we estimate the variance using a 1×3 kernel, which allows to capture the vertical variance in the temperature map in each row. Afterwards, we apply an edge-sensitive bilateral filter [16] with a 7×7 kernel to smooth out the variance map while preserving the edges of the separating regions.

Due to perspective and varying distance between the modules and the infrared camera, module rows are of varying size, shape, and temperature. Thus, we process each row separately. We equalize the histogram of the variance map corresponding to each row and threshold the variance map to zero at 20 % of the map range. This amplifies the separating regions, which are processed using a sequence of morphological dilation, closing, and erosion using large horizontal 7×1 kernels.

By computing the average row width, we fit a rectangle to high variance regions that separate the modules, and estimate their orientation using ellipse fitting. This approach, however, may result in inaccurate orientation estimates if the separating regions are very small. Therefore, if the orientation of the separating region deviates over 50 % from the orientation of its immediate neighbors, we determine the orientation by linearly interpolating it between the orientation of neighboring separating regions. In case the separator is the first or the last in the module row, we use only the separating region of its immediate neighbor to correct the orientation.

4.6. Correction of Module Dimensions

The segmentation results so far provide a robust estimation of individual module regions in the temperature map. However, in certain cases modules at the start or the end of the row can only be detected partially. This problem occurs if the temperature decreases towards the edge of the module row resulting in portions of the module to be cut off, because these regions fall under the previously determined segmentation threshold. As the module height is typically unaffected by this problem, we correct only the module width.

To correct the module width, we compute the average module width in each row, and extend the modules at the beginning and the end of the row to the determined width.

4.7. Preparing Modules for Analysis

Once the dimensions of each segmented photovoltaic module are corrected, we extract the temperature map regions belonging to individual modules, and unwarp their perspective using the corner points. This step creates a bird's eye view of the modules, which are slightly tilted in our dataset, finally allowing to assess the modules for defects.

5. Analysis of Photovoltaic Modules Using Infrared Thermography

Within a photovoltaic module, a wide range of defects can occur. Many of these defects cause overheating either

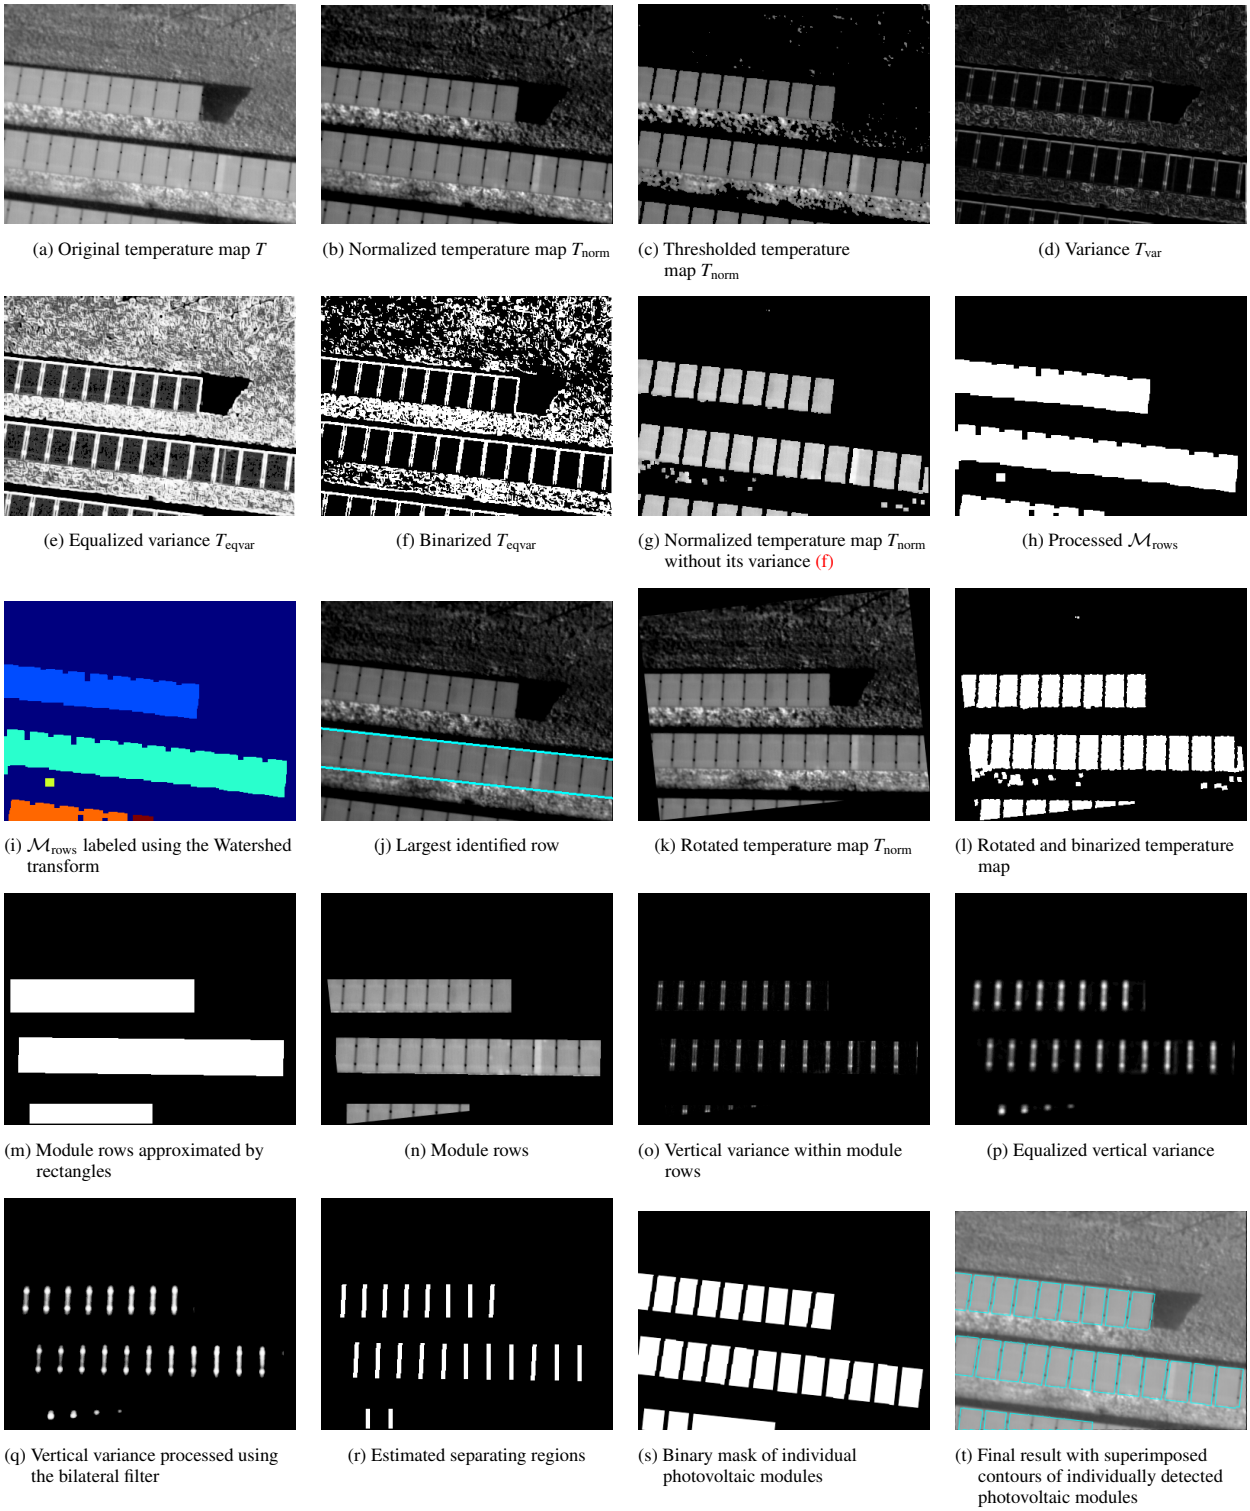

(a) Original temperature map T

(b) Normalized temperature map T_{norm}

(c) Thresholded temperature map T_{norm}

(d) Variance T_{var}

(e) Equalized variance T_{eqvar}

(f) Binarized T_{eqvar}

(g) Normalized temperature map T_{norm} without its variance (f)

(h) Processed $\mathcal{M}_{\text{rows}}$

(i) $\mathcal{M}_{\text{rows}}$ labeled using the Watershed transform

(j) Largest identified row

(k) Rotated temperature map T_{norm}

(l) Rotated and binarized temperature map

(m) Module rows approximated by rectangles

(n) Module rows

(o) Vertical variance within module rows

(p) Equalized vertical variance

(q) Vertical variance processed using the bilateral filter

(r) Estimated separating regions

(s) Binary mask of individual photovoltaic modules

(t) Final result with superimposed contours of individually detected photovoltaic modules

Figure 3. Intermediate and final results of our photovoltaic module detection approach

of the entire module, single substrings, small solar cells, or even tiny cell parts. Under infrared thermography, such defects can produce various patterns. However, not all patterns are a clear sign of module malfunction. For instance, photovoltaic modules may exhibit regular patterns which are simply characteristic to the photovoltaic module types being used.

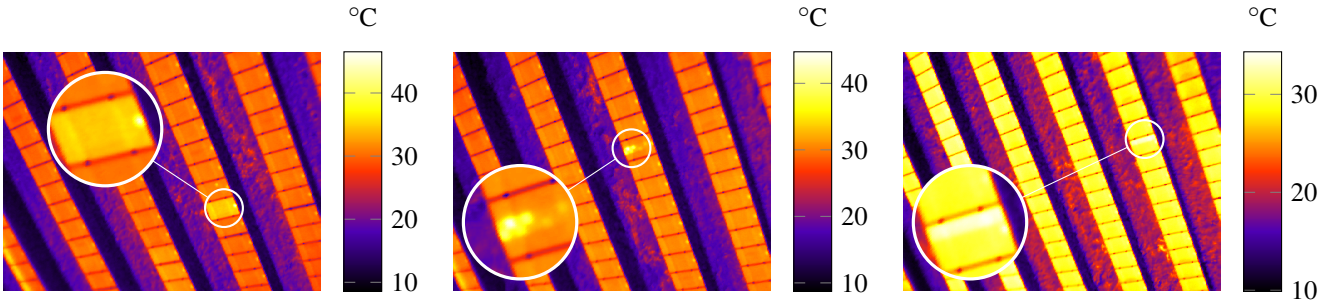

Figure 4. Examples of defects in photovoltaic modules which we classify using our approach: overheated modules (left), hot spots (center), and overheated substrings (right).

In our pipeline, we analyze all detected photovoltaic modules within an infrared image at once using statistical tests. We compute four sets of features, which are tested for outliers indicating a temperature abnormality within each photovoltaic module. A cascade of two statistical tests, Grubbs' [17] and Dixon's Q test [18], is used to verify the mutual outcome. The test cascade increases the classification accuracy, since the tests complement each other.

The four feature sets 1) module medians, 2) grid cell medians, 3) histogram skewness, and 4) vertical projections, which correspond to the samples being statistically tested for outliers, address different types of defects in photovoltaic modules.

5.1. Detection of Defective Photovoltaic Modules

We examine four feature sets to detect and classify defects in photovoltaic modules.

Module Medians The feature set containing median values of each photovoltaic module, which allows to detect completely *overheated modules*.

Grid Cell Medians The feature set of median values per grid cell in a module, which can be used to identify *hot spots*.

Histogram Skewness The feature set of skewness factors of a histogram of median values per grid cell. This feature set is suitable for detecting *overheated substrings*.

Vertical Projections capture the irradiance profile while considering module's topology. Vertical projections complement the histogram skewness feature set allowing to identify *overheated substrings*.

Fig. 4 illustrates the different types of defects.

5.1.1 Overheated Modules

Overheated photovoltaic modules emit a considerably higher temperature than those operating at a normal one. Deviations from median temperature of several photovoltaic modules within an infrared image are an indicator for such abnormalities.

Due to overheating, the median temperature is typically around 2 °C higher than the normal operating temperature. However, overheating also heavily depends on the weather conditions, which can result in temperature fluctuations of 1 °C to 10 °C. Moreover, since photovoltaic modules are connected to each other, overheating in one module may cause overheating in another one. Using the median temperature we are able to detect such defects.

To identify the overheated modules, we compute the median temperature $\{m_k\}_{k=1,\dots,\ell}$ for every photovoltaic module k. Here, we denote $\ell \in \mathbb{N}$ as the number of modules within an infrared image. Ideally, the medians should be normally distributed. An outlier will thus indicate a temperature abnormality. By defining the medians as a set of observations or equivalently a set of features $\mathcal{F}_1$, we subsequently apply a cascade of statistical tests to detect the outliers.

5.1.2 Hot Spots

Modern photovoltaic modules often consist of 6×10 solar cells connected to three substrings. As a consequence, overheating can also affect just certain areas of a photovoltaic module instead of the entire module surface. In this regard, imaged temperature values within a photovoltaic module are not independent from each other, and thus can be classified in batch. Therefore, we divide the modules in a $N \times M$ grid of cells, and compute the median temperature ${}^{j}m_k \in \mathbb{R}$ within the corresponding temperature map regions of the grid cell $j \in \{1, \dots, NM\}$ associated with the module k.

Dividing the temperature map of a module into a grid of cells provides two major benefits. First, we effectively reduce the search space and as a consequence the computational complexity. Secondly, the grid division allows to minimize the influence of regular patterns within a photovoltaic module which can negatively impact the classification performance. A division of 9×10 cells proved to produce most accurate classification results for our dataset.

Similar to module medians, grid cell medians are tested for outliers using the statistical test cascade in order to detect overheated areas.

5.1.3 Overheated Substrings

Modern photovoltaic modules with 6×10 solar cells consist of three substrings with 20 solar cells each: left, middle, and right section. When the series connection of the solar cells belonging to one bypass diode is short-circuited, the substring does not contribute to the power generation of the photovoltaic module. The short circuit can be caused by either a defective, short-circuited bypass diode, or by an intact bypass diode overriding defective or shaded solar cells that act as a resistor in the series connection.

In the case of an inactive substring, the incident irradiation energy is not led away and heats the affected solar cells. These in turn warm up the overlying front side glass area of the photovoltaic module due to heat conduction. As a result, an inactive substring causes overheated vertical patterns along module's long edge. We identify such patterns using histogram skewness and vertical projections of grid cell medians.

Histogram Skewness To identify substring defects, we first build a reference histogram of grid cell medians ${}^{j}m_{\hat{k}}$ of every photovoltaic module $\hat{k} \in \{1,\dots,k-1,k+1,\dots,\ell\}$ in the infrared image excluding module k. Meanwhile, we treat the grid cell medians ${}^{j}m_{k}$ of each module k as the bins of a second histogram and compute its skewness v_k with respect to the reference histogram as

$$v_k \coloneqq \frac{1}{NM} \sum_{j=1}^{NM} \left(\frac{{}^{j}m_k - \bar{m}}{s} \right)^3 , \tag{6}$$

where $\bar{m}$ is the sample mean, and s the sample standard deviation estimated from the reference histogram. The skewness factors $\{v_k\}_{k=1,\dots,\ell}$ correspond to the third feature set $\mathcal{F}_3$, which is tested for outliers using the cascade of statistical tests.

Vertical Projections Histogram skewness is not always a reliable feature for detection of overheated substrings. In infrared images, where the temperature across the detected photovoltaic modules fluctuates heavily, or the defective substrings appear very narrow due to strong perspective, histogram skewness may not produce distinctive features. Based on the work of Fang *et al.* [19] we therefore complement histogram skewness by the vertical projection $\rho_k^{(i)} \in \mathbb{R}$ of grid cell medians to detect defective substrings. The vertical projections are the means along each grid column i of the module k. For a $N \times M$ grid of cells, we compute M vertical projections per each detected module. To eventually identify defective substrings, the observations $\rho_k^{(i)}$ are tested for outliers by the cascade of statistical tests.

5.2. Classification of Defects Using Statistical Tests

To identify defective photovoltaic modules, we apply two statistical outlier tests to the feature sets $\{\mathcal{F}_i\}_{i=1,\dots,4}$, where

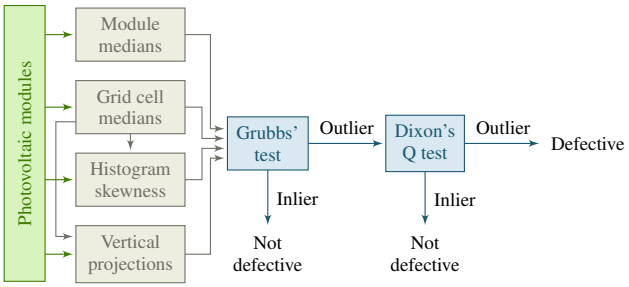

Figure 5. The cascade of outlier tests, which we use to detect various defects within a photovoltaic module using four feature sets.

$\mathcal{F}_i \coloneqq \{f_1, f_2, \dots, f_n\}$. By applying the statistical tests, we examine the distribution of $\mathcal{F}_i$ for values that do not belong to the normal distribution.

We use a cascade of two outlier tests: 1) Grubbs' test for outliers followed by 2) Dixon's Q test, whereat Dixon's Q test is used to verify the outcome of Grubbs' test (see Fig. 5). Our null hypothesis states that there are no significant outliers in the feature set, or in other words: there are no temperature abnormalities within each detected photovoltaic module. In case the null hypothesis is rejected by both tests, the photovoltaic module is regarded as being defective.

5.2.1 Grubbs' Test

Grubbs' test for outliers is an iteratively performed test, which can detect one outlier at a time. The test is applied to the feature set $\mathcal{F}_i \coloneqq \{f_1, f_2, \dots, f_n\}$ with its observations $f_1 \leq f_2 \leq \dots \leq f_n$ arranged in ascending order. Detected outliers are removed from the set until no outliers are detected. The test statistic is given by

$$G_{\max} \coloneqq \frac{f_n - \mu}{\sigma} , \tag{7}$$

where $f_n \coloneqq \max \mathcal{F}$ is the largest value in the set. The critical value of the Grubbs' test is given by

$$G_{\text{crit}} \coloneqq \frac{(n-1)\, t_{\alpha/n,n-2}}{\sqrt{n\left(n-2+t^2_{\alpha/n,n-2}\right)}} , \tag{8}$$

where $t_{\alpha/n,n-2}$ is the critical value of Student's t-distribution at the significance level α and $\nu \coloneqq n-2$ degrees of freedom. The null hypothesis is rejected, if $G_{\max} > G_{\text{crit}}$, *i.e.*, the maximum value does not belong to the normal distribution and thus points to a defective photovoltaic module.

We perform a total of $m \coloneqq n-2$ tests, where the maximum number of iterations is limited by the minimum number of degrees of freedom ν of the t-distribution, which is 1.

To increase the classification accuracy, we introduced two additional criteria that support Grubbs' test decision. First, the difference between the outlier (corresponding to

the maximum value in $\mathcal{F}_i$) and sample mean must be larger than 1 °C. Secondly, the absolute difference between the two largest values in the feature set must be at least 1.5 °C. Both criteria were determined empirically and in accordance with the typical temperature fluctuations of photovoltaic modules used in our experimental setup.

5.2.2 Dixon's Q Test

Similar to Grubbs' test, Dixon's Q test is applied to observations $f_1 \leq f_2 \leq \cdots \leq f_n$ of the feature set $\mathcal{F}_i \coloneqq \{f_1, f_2, \ldots, f_n\}$. The test statistic is given by

$$Q \coloneqq \frac{f_{n-1} - f_n}{f_1 - f_n} \quad , \tag{9}$$

where $f_n \coloneqq \max \mathcal{F}$ is the largest and f_{n-1} the second largest value in $\mathcal{F}$, respectively. $f_1 \coloneqq \min \mathcal{F}$ is the smallest value in the set. We reject the null hypothesis, if $Q > Q_{\text{crit}}$, where Q_{crit} is the reference value corresponding to the cardinality of $\mathcal{F}_i$ and the significance level α.

A major disadvantage of Dixon's Q test is its limitation to observation sets not larger than $n \leq 30$, whereas Grubbs' test for outliers can be applied to arbitrary number of observations. To deal with this deficiency, we apply Dixon's Q test only to a subset $\mathcal{F}_i' \subset \mathcal{F}_i$ of most significant observations, *i.e.*, the largest values in $\mathcal{F}_i$. We limit the number of observations in $\mathcal{F}_i'$ to $|\mathcal{F}_i'| = 25$, which is sufficient to classify defective modules accurately.

6. Experimental Results

The performance of the photovoltaic module detection was evaluated on 24 infrared images. The test images were hand labeled. The evaluation was performed by comparing the final segmentation mask produced by the algorithm with the hand-labeled reference mask.

The sensitivity of the detection algorithm on the test images reached 97.66 % and the specificity 95.54 %, which indicates an above average performance of the algorithm at identifying true positives and true negatives. The F_1 score is a measure of accuracy incorporating both the precision and recall and can be used as a single performance test for positive classifications. F_1 score is defined as a weighted average between precision and recall which reached 92.76 % for our dataset. We additionally computed the Matthews correlation coefficient (MCC), which yields balanced results even if the true and negative classes are of different sizes. This was of interest for infrared images which contained few photovoltaic modules and thus only few true positives. An MCC of 89.55 % indicated a high correlation between ground truth and algorithm's predictions. Table 1(a) summarizes these results.

The classification of defects in photovoltaic modules was evaluated on 37 infrared images containing 1544 modules which were labeled by hand. For Grubbs' test we used a significance level of $\alpha_G \coloneqq 0.01$ to classify overheated modules. For remaining features, a significance level of $\alpha_G \coloneqq 0.1$ provided the best performance. Concerning Dixon's Q test, a significance level of $\alpha_Q \coloneqq 0.01$ was used for all the features. This parametrization resulted in an F_1 score of 93.88 % indicating an above average classification performance. The results of the classification component are summarized in Table 1(b).

Measure	p
Sensitivity	97.66 %
Specificity	95.54 %
Precision	88.34 %
Recall	97.66 %
Accuracy	96.19 %
MCC	89.55 %
F_1 score	92.76 %

(a) Detection of photovoltaic modules

Measure	p
Sensitivity	97.83 %
Specificity	99.67 %
Precision	90.20 %
Recall	97.83 %
Type I error	0.27 %
Type II error	2.17 %
F_1 score	93.88 %

(b) Classification of defects in photovoltaic modules

Table 1. The performance of our photovoltaic module detection and analysis pipeline

7. Conclusions

The proposed image processing pipeline for the automatic evaluation of aerial infrared imagery of photovoltaic power plants provides a solid basis for future automatic monitoring and early detection systems. The two components for detection of photovoltaic modules and classification of defects in modules are fundamental to any future analysis software.

As shown in Fig. 3(t), our image processing pipeline detects individual photovoltaic modules within an aerial infrared image reliably. Also, major temperature abnormalities which are often caused by module defects are detected and classified accurately into three main groups: overheated modules, hot spots, and overheated substrings.

In combination with further data on the installed photovoltaic modules and information about the spatial and electrical layout of the photovoltaic power plant, the detection of photovoltaic modules as well as the interpretation of identified defects can be optimized.

Acknowledgments

This work is part of the "Smart Grid Solar" project co-financed by the European Union through the European Regional Development Fund, and by the Free State of Bavaria. This work was partly supported by the Research Training Group 1773 "Heterogeneous Image Systems" funded by the German Research Foundation (DFG). The authors would like to thank IBC Solar AG for the possibility to use several photovoltaic power plants as test objects for the development of pattern recognition algorithms.

References

[1] Fraunhofer-Institut für Solare Energiesysteme ISE, "Aktuelle Fakten zur Photovoltaik in Deutschland," 2015, (German). 1

[2] E. Molenbroek, D. W. Waddington, and K. A. Emery, "Hot spot susceptibility and testing of PV modules," in *IEEE Photovoltaic Specialists Conference*, 1991, pp. 547–552. 1

[3] W. Herrmann, W. Wiesner, and W. Vaassen, "Hot spot investigations on PV modules – new concepts concepts for a test standard and consequences for with respect to bypass diodes," in *IEEE Photovoltaic Specialists Conference*, 1997, pp. 1129–1132. 1

[4] J. Wohlgemuth and W. Herrmann, "Hot spot tests for crystalline silicon modules," in *IEEE Photovoltaic Specialists Conference*, 2005, pp. 1062–1065. 1

[5] C. Buerhop, D. Schlegel, M. Niess, C. Vodermayer, R. Weißmann, and C. J. Brabec, "Reliability of IR-imaging of PV-plants under operating conditions," *Solar Energy Materials and Solar Cells*, vol. 107, pp. 154–164, 2012. 1

[6] C. Buerhop, R. Weißmann, H. Scheuerpflug, R. Auer, and C. J. Brabec, "Quality Control of PV-Modules in the Field Using a Remote-Controlled Drone with an Infrared Camera," in *Proceedings of the 27th European Photovoltaic Solar Energy Conference and Exhibition (EU PVSEC 2012)*, 2012, pp. 3370–3373. 1

[7] C. Buerhop and H. Scheuerpflug, "Field inspection of PV-modules using aerial, drone-mounted thermography," in *Proceedings of the 29th European Photovoltaic Solar Energy Conference and Exhibition (EU PVSEC 2014)*, 2014, pp. 2975–2979. 1

[8] J. R. M. de Dios and A. Ollero, "Automatic detection of windows thermal heat losses in buildings using UAVs," in *World Automation Congress*, Jul. 2006, pp. 1–6. 2

[9] J. A. Tsanakas, D. Chrysostomou, P. N. Botsaris, and A. Gasteratos, "Fault diagnosis of photovoltaic modules through image processing and Canny edge detection on field thermographic measurements," *International Journal of Sustainable Energy*, vol. 34, no. 6, pp. 351–372, Jul. 2015. 2

[10] G. Leotta, P. M. Pugliatti, A. D. Stefano, F. Aleo, and F. Bizzarri, "Post processing technique for thermo-graphic images provided by drone inspections," in *Proceedings of the 31st European Photovoltaic Solar Energy Conference and Exhibition (EU PVSEC 2015)*, 2015, pp. 1799–1803. 2

[11] DaVinci Copters. (2015) ScaraBot construction kit. [Online]. Available: http://davinci-copters.com/scarabot-construction-kit 2

[12] ZAE Bayern, "Abschlussbericht der Machbarkeitsstudie zur Überprüfung der Qualität von Photovoltaik-Modulen mittels Infrarot-Aufnahmen," 2007, (German). 2

[13] C. Rother, V. Kolmogorov, and A. Blake, "'GrabCut': Interactive foreground extraction using iterated graph cuts," in *ACM SIGGRAPH*, ser. SIGGRAPH '04. New York, NY, USA: ACM, 2004, pp. 309–314. 3

[14] F. Meyer and S. Beucher, "Morphological segmentation," *Journal of Visual Communication and Image Representation*, vol. 1, no. 1, pp. 21–46, Sep. 1990. 4

[15] S. Dotenco, F. Gallwitz, and E. Angelopoulou, "Autonomous approach and landing for a low-cost quadrotor using monocular cameras," in *Computer Vision – ECCV 2014 Workshops*, ser. Lecture Notes in Computer Science, L. Agapito, M. M. Bronstein, and C. Rother, Eds. Springer International Publishing, 2015, vol. 8925, pp. 209–222. 4

[16] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proceedings of the 6th International Conference on Computer Vision*. Narosa Publishing House, 1998, pp. 839–846. 4

[17] F. E. Grubbs, "Sample criteria for testing outlying observations," *The Annals of Mathematical Statistics*, vol. 21, no. 1, pp. 27–58, Mar. 1950. 6

[18] R. B. Dean and W. J. Dixon, "Simplified statistics for small numbers of observations," *Analytical Chemistry*, vol. 23, no. 4, pp. 636–638, 1951. 6

[19] Y. Fang, K. Yamada, Y. Ninomiya, B. Horn, and I. Masaki, "Comparison between infrared-image-based and visible-image-based approaches for pedestrian detection," in *Proceedings of the IEEE Intelligent Vehicles Symposium*. IEEE, 2003, pp. 505–510. 7