

Continuous Short-term Irradiance Forecasts using Sky Images

David Bernecker, Christian Riess, Elli Angelopoulou, Joachim Hornegger

Pattern Recognition Lab, Department of Computer Science, Friedrich-Alexander-Universität Erlangen-Nürnberg, Martensstr. 3, 91058 Erlangen, Germany

Abstract

We present a system for forecasting occlusions of the sun and the expected Global Horizontal Irradiance (GHI) for solar power plants. Our system uses non-rigid registration for detecting cloud motion and a Kalman filter to establish continuous forecasts for up to 10 minutes. The optimal parameters of the system are determined through the use of the binary classification metrics Precision, Recall and F_2 Score while evaluating the forecasting of occlusions. The Kalman filter and the use of a dense motion field instead of a global cloud speed prove to be key elements of the forecasting pipeline: by incorporating information from previous forecasts into the current one, a Kalman Filtering facilitates forecasting times below 3 minutes and the dense motion field enhances the accuracy of our forecasts. Our evaluation of the proposed approach on 15 days of real world data collected in Kitzingen, Bavaria, Germany, produced a mean RMSE for forecasting GHI of $(164 \pm 9) \text{W m}^{-2}$.

Keywords: Solar forecasting; Sky imaging; Global Horizontal Irradiance (GHI); Intra-hour forecast

1. Introduction

An important factor for the wider adoption of photovoltaics (PV) as a renewable energy source on a larger scale, is stable power production. This is a challenging task since sun-occlusion effects by clouds can cause significant variability in the power output of solar power plants. Unaccounted changes in the future production of the power plant may either lead to an overproduction of power, or it may require the compensation of underproduction by additional power sources at short notice. Both of these events decrease the cost effectiveness of the power plant and should, therefore, be avoided. One way to improve the profitability and reliability of solar power plants is by forecasting the future power production, and therefore allowing for an earlier adjustment to the present conditions. We propose a new methodology for estimating upcoming sun occlusions for forecasts of up to 10 minutes (depending on the current cloud speeds) by analyzing sky-images acquired from the ground.

Our forecasts are quasi-continuous and describe the probability of the sun being occluded in a step size as small as 2 seconds. From such forecasts, one can calculate the periods of time in the future when the sun is going to be occluded. The main contributions of this work are:

- The generation of continuous forecasts for time intervals of up to 10 minutes, depending on the current cloud motion speed, instead of individual time-instance forecasting.
- The application of non-rigid registration to the task of estimating the cloud speed. Instead of a global

cloud speed, we determine the cloud speed individually for different distances to the sun, thus the proposed method gains flexibility to account for deformations of the clouds.

- The use of a Kalman filter (Kalman, 1960) to improve the accuracy of the forecast by incorporating knowledge from previous forecasts. The result is a continuous forecast also for forecasting times below 3 minutes.
- The quantitative evaluation of occlusion periods on real world data using Precision, Recall and F_2 Score as metrics, as well as the evaluation of our method using the Root Mean Square Error of predicted irradiance and the forecast skill. Our proposed method outperforms the persistence baseline even for short forecasting times.

The remainder of this paper is structured as follows. In section 2 we present the current state of the art for forecasting PV power production. Section 3 contains a detailed description of our method and in section 4 the conducted experiments and their evaluation are presented.

2. State of the Art

The approach used for predicting the solar irradiance depends on the desired time scale of the forecast. For longer forecast times, in the range of 6 hours to several days, numerical weather models can establish predictions for the expected irradiance for whole regions of interest (Lorenz et al., 2009). For predictions over a couple of hours, the

cloud velocity can be extracted from satellite images. By extrapolating the trajectories of clouds in these images, the effects of occlusions can be predicted (Kidd et al., 2009). For an even higher temporal resolution, reaching down to several minutes, sky images recorded from the ground (Chow et al., 2011), or time-series models that rely on the knowledge of historical data (Reikard, 2009) can be employed. Our method only uses ground-based sky images. We will, therefore, concentrate on the comparison to similar methods in this brief overview of previous work. For more detailed information on the other types of methods, one may refer to the recent review by Inman et al. (2013).

Chow et al. (2011) were the first to establish intra-hour forecasts for solar irradiance using a ground-based camera. For recording the image data a commercially available Total Sky Imager from Yankee Environmental Systems was used. Cloud Segmentation was done by employing a threshold on the ratio of the red/blue channels of the images and by also comparing the sky images to a clear-sky model. Actual predictions are done by calculating the cloud speed using the cross-correlation of cloud patches, and then moving the cloud mask according to the determined cloud speed. The whole approach was evaluated on a four day dataset. However, the method and the used equipment have some drawbacks. The resolution of the acquired images is low, and the default automatic white-balancing can lead to problems in cloud segmentation. Furthermore, images are recorded in the JPEG format, which induces losses in image quality during compression. The method for assessing the cloud speed is not robust against non-rigid deformations of the clouds, and the quantitative part of the evaluation is restricted to how well the forecasted cloud map corresponds to the actual one. No irradiance data is used in the quantitative part of the evaluation.

An improvement in the cloud motion registration was proposed by Marquez and Coimbra (2013), who use particle image velocimetry (PIV), a technique commonly used in the experimental study of fluid dynamics. The cloud segmentation was also further improved by using an adaptive thresholding method, and by partitioning the sky image into several regions with different thresholds. For evaluation, Marquez and Coimbra used the Root-Mean-Squared error (RMSE) between the forecasted Direct Normal Irradiance (DNI) and the measured one.

Two publications built up on this work: Chu et al. (2013) further improved the forecasts by training an artificial neural network (ANN), while Quesada-Ruiz et al. (2014) replace the PIV approach by a sector based motion detection method. The ANN required as input the expected cloudiness calculated by the method of Marquez and Coimbra and the last few measured irradiance values. The result of the ANN was a 5 or 10 minutes irradiance forecast. The latter work by Quesada-Ruiz et al. (2014) divided the sky into sectors, and determined the direction of cloud motion by analyzing the changes of cloudiness




Figure 1: Example sky image.

in these sectors. Using the RMSE between predicted and measured DNI as a measure, the evaluation of both works showed an improvement over the previous work by Marquez and Coimbra.

Besides using ground-based images, there also exist approaches that only rely on a network of irradiance sensors. Lonij et al. (2013) for instance, established intra-hour forecasts using a network of 80 rooftop PV systems spread over a large area. For large but isolated solar power plants the cloud motion vectors could also be calculated using reference cells, as was shown by Bosch and Kleissl (2013). Both of these approaches, however, require additional hardware that has to be installed in a large region around the power plant. In comparison, the Sky Imager only requires a single camera that is usually located at the center of the power plant.

Further work has also been done on the segmentation and classification of clouds. Kazantzidis and Tzoumanikas (2012) augmented the cloud segmentation based on the red/blue ratio by also incorporating the green channel. Furthermore, they used texture features and a k -Nearest-Neighbor classifier to classify the type of cloud seen in an image. However, the classification is done for the entire image, so that sky conditions where more than one type of cloud is present lead to an ambiguous result.

3. Forecasting Pipeline

We will first give an overview of the overall system architecture. The key components of our approach are then presented in greater detail in the remainder of this section.

3.1. Overview

The whole pipeline for establishing forecasts is presented in Fig. 2. Our starting point is the acquisition of the sky images (cf. Fig. 1). The clouds are then segmented and the




Figure 2: Pipeline of the algorithm for establishing a forecast.

cloud speed (in pixels per second) is determined. These two steps can be performed in parallel. For segmenting the clouds, we compare a threshold of the red/blue ratio as used by [Chow et al. \(2011\)](#) and an approach based on a trained classifier that takes into account the structure of the clouds. In contrast to previous publications on the determination of cloud speed, we use a method that can cope with non-rigid deformations of the clouds between successive frames, namely the non-rigid registration method by [Thirion \(1998\)](#). Afterwards, the results of the segmentation and the cloud motion field are combined to establish a preliminary forecast. Using a Kalman Filter this preliminary forecast is combined with previous forecasts to calculate the final forecast. The output of our forecasting algorithm is the expected occlusion probability of the sun for the next 10 minutes which can be used to predict the global horizontal irradiance (GHI). **Such a forecast is issued every 5 seconds.** In the subsequent sections, we describe each step of the pipeline in greater detail.




Figure 3: Example displacement vectors calculated with the "demons" algorithm (only every 20th displacement vector is shown).

3.2. Data acquisition & preprocessing

The sky images are acquired with a five megapixel C-mount camera equipped with a fish eye lens. The camera-lens combination provides a large field of view of 182° . The size of the images is 2592×1944 pixels. The projected image occupies a circular region with a diameter of 1944 pixels (cf. Fig. 1). Images with a fixed exposure time of $136 \mu\text{s}$ and preset white balance setting are acquired every 5 seconds. The parameters for white balancing are set manually and are kept constant over the whole evaluation period. This is necessary since the cloud segmentation (cf. section 3.4) relies on the ratio of different color channels. The camera is installed next to several pyranometers which record the global horizontal irradiance (GHI) in 5-second intervals.

The mapping function of the fisheye lens is assumed to be an equidistant projection (i.e., angular distances are maintained by the projection). Given this mapping function and the focal length of the lens, the images can easily be rectified. Furthermore the orientation of the camera is corrected, so that the center point of the image corresponds to the zenith in the sky and the bottom edge of the image faces south.

3.3. Cloud-speed estimation

Thirion's "demons" algorithm ([Thirion, 1998](#)) is an image registration approach that is commonly used in medical imaging. In this work we use it to determine the locally dominant movement direction and speed of the clouds over




Figure 4: Example distribution of displacements vectors between two successive sky images. The mode of the distribution (black cross) is assumed to correspond to the global cloud displacement.

the visible sky. Please note that cloud speed in this context means the speed of the clouds in *pixels per minute*. Since the height of the clouds is unknown, a conversion to the actual velocity in meters per second is not possible. Thirion’s algorithm maps two consecutive cloud images in a non-rigid manner: it estimates not only a single global translation or rotation, but instead a deformation vector field which incorporates small turbulences and locally varying motion.

The algorithm itself is based on the brightness constancy assumption

$$I_t(\mathbf{p}) = I_{t+\Delta t}(\mathbf{p} + \mathbf{d}) \quad (1)$$

that states that image intensities of moving objects remain constant. The recorded intensity values can only change location \mathbf{p} . The deformation field $\mathbf{d} = (d_x, d_y)^T$ that maps two consecutive images (I_t, I_{t+1}) onto each other, is determined in an iterative way by minimizing equation (1) for all pixels in the image (cf. [Appendix A](#)).

The result of the registration method is a deformation field that assigns a displacement vector to each pixel. To extract the principal direction of the cloud motion, some post-processing of this vector field is necessary. Since all clear sky pixels have a displacement of zero, taking the average of the displacement field is not sufficient. Instead, we search for the mode of the distribution of displacement vectors (i. e., the displacement value with the highest number of occurrences) after removing all displacements close to zero. A typical distribution of displacement vectors is shown in [Fig. 4](#) with the detected mode marked by a black cross.

To reduce the effect of fluctuations in the measured principal direction, a secondary Kalman filter (cf. [Appendix B](#)) with an identity matrix as a state transition matrix is used inside the pipeline step that determines the cloud speed. By tracking the principal direction over time we obtain a more accurate estimate. Both the principal direction and the pixel by pixel displacement field is used in our forecast.

In contrast to previously used methods like block-wise cross-correlation, the non-rigid registration also captures deformations of the clouds. The differences between using a global motion vector and the displacement field are illustrated in [Fig. 5](#). For an image pair the global motion was determined using the cross-correlation method from [Chow et al. \(2011\)](#), and the deformation field was determined using the demons algorithm. The resulting motion vector and deformation field were used to translate or warp one of the input images to match the second image. Shown in [Fig. 5](#) are the absolute differences between the transformed and the target image. Brighter pixels indicate larger registration errors. As can be seen, the demons algorithm captures the deformations of the clouds more precisely.

3.4. Cloud segmentation

We investigate two approaches for the identification of cloud pixels. The basic method exploits the high contrast

between cloud and sky pixels when looking at the ratio of their red and blue channels ([Chow et al., 2011](#)). In general, cloud pixels should have a value close to one, while sky pixels should exhibit lower values. **The decision between cloud and sky pixels is done by comparing the red/blue ratio to a threshold Θ_{rb} . For a ratio larger than Θ_{rb} a pixel is identified as a cloud pixel, and for values smaller than the threshold Θ_{rb} as a sky pixel.** While computationally very efficient, this method becomes unreliable for thick clouds that do not let any direct sunlight through. These clouds are only illuminated by light reflected from the Earth’s surface and have a low red/blue ratio that is very similar to sky pixels. This leads to an underdetection of cloud pixels and consequently to an overestimation of free sky.

The second method augments the color information with additional texture features. However, cloud texture is subject to high variability, which makes it hard to explicitly describe. To address this issue, we employ a recently proposed representation-learning algorithm ([Bernecker et al., 2013](#)). Using large amounts of image patches showing either clouds, sky or a mixture of both, texture filters, which specifically describe cloud texture, can be learned using a Restricted Boltzmann Machine, a form of artificial neural network. To detect clouds, the sky image is convolved with these learned filters, and for each pixel the filter responses and the red/blue ratio are stacked into a vector. This feature vector is then classified using a Random Forest classifier, i. e., each pixel is individually classified as cloud or sky pixel. The training of the classifier in this approach consists of two stages. First, the filters corresponding to cloud texture are learned from 160 000 image patches that are randomly extracted from sky images. In the second stage the Random Forest is trained. This is done using 350 000 manually labeled pixels from several sky images. A cross-validation scheme is used to determine the optimal parameters of the classifier. Another 150 000 sample pixels, again labeled by hand, are then used to evaluate the approach.

Compared to the work by [Kazantzidis and Tzoumanikas \(2012\)](#), this method classifies clouds on a pixel-by-pixel basis, and not the whole image at once. Furthermore, only the presence of a cloud is classified, but not its type. An extension to classify different types of clouds is, however, possible when sky images with labeled cloud types are available.

Our experiments, however, have shown that although the recognition rate of clouds is better with the learning-based method, the required computation time is too high to be incorporated into our pipeline. Therefore, we ended up using the simple thresholding-based method in our evaluation.

3.5. Cloud forecasting

Our forecasting algorithm takes the segmented clouds and the cloud speed determined from the two previous steps as an input. The resulting forecast is the estimated

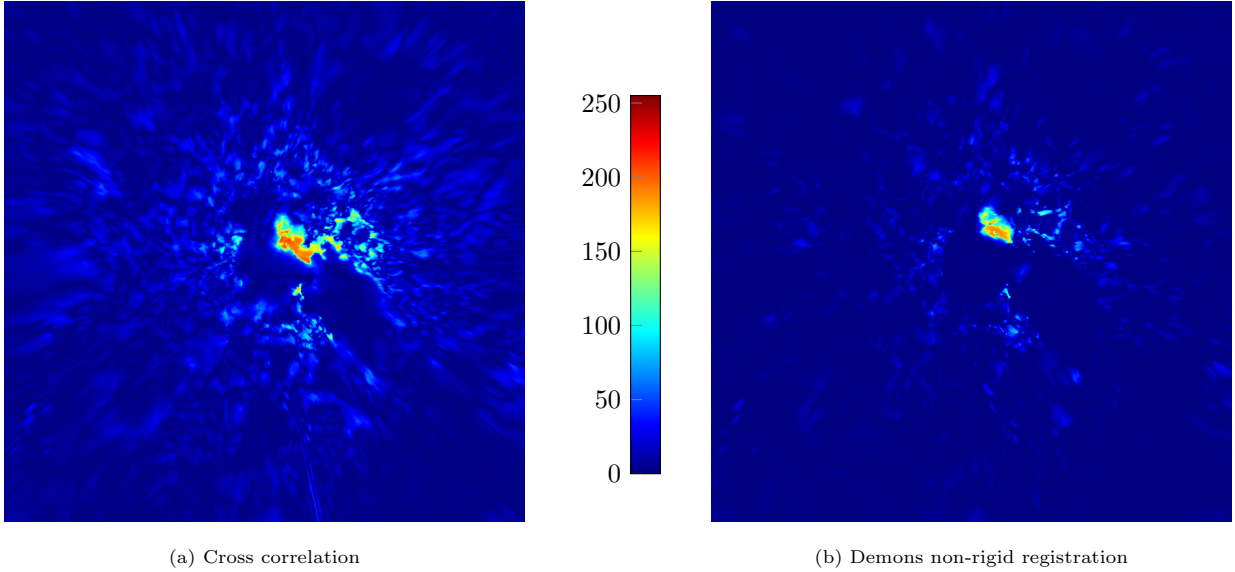


Figure 5: Absolute differences between two registered images. **The color indicates the value of the difference.** The proposed approach using Thirions demons registration method better accounts for local cloud deformations.

probability of occlusion at the position of the sun during the next 10 minutes.

For establishing a forecast, we use only the part of the sky with clouds moving towards the position of the sun. More specifically, we use a circular sector of the sky (cf. Fig. 6a), whose tip is at the position of the sun. The opening angle γ_{\max} of the circular sector is one of the parameters of our forecast that are determined in a sensitivity analysis (cf. section 4.3). The height and azimuth of the sun are easily calculated using the geographical position of the camera and the time of day. They can be converted to image coordinates using the known mapping function of the lens. The central angle is oriented opposite to the principal direction of cloud movement that was obtained by finding the mode of the motion vectors, so that within the cropped area, clouds are moving towards the sun. The radius of the circle depends on the estimated cloud speed, so that clouds that might intersect the sun in the next 10 minutes are within the segment.

Using a polar transformation, the circular sector is transformed to a rectangular image, where the x-axis corresponds to the distance to the sun in pixels, and the y-axis to the angle γ in the cone (cf. Fig. 6b). The *cloud chart* (cf. Fig. 6c) is determined by applying the same coordinate transformation on the segmented image. Note that during the polar transform the area close to the sun is sampled more densely than the area farther away.

We then calculate the weighted sum of each column in the cloud chart that provides the occlusion probability for the next few minutes (cf. Fig. 6d). As a weighting function we use a Gaussian with a high standard deviation (chosen so that the value of the Gauss function at the top and bottom of the cloud chart (cf. Fig. 6c) is half of the value at the center of the column). We have chosen this weighting since the probability of occluding the sun is higher for

clouds on the center line of the cloud chart. Finally, the conversion from the distance to the sun to the estimated time until the sun is reached is done for each column separately by averaging over the corresponding motion vectors where a cloud was identified in the cloud chart. This results in the basic forecast. Since not for all distances an arrival time can be calculated (e.g. for columns only containing clear sky pixels with zero displacement vectors) this preliminary forecast can be sparse.

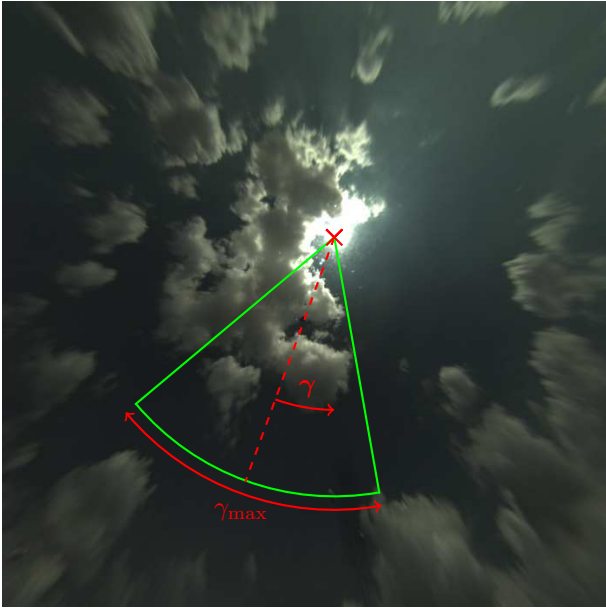
Due to saturated pixels close to the sun, the sparsity of the preliminary forecast, as well as inaccuracies introduced in the previous steps (e.g. cloud segmentation), the forecast obtained from one cloud chart is not yet a reliable predictor of sun occlusion. Hence, we use a second Kalman filter to exploit information from past measurements and get more accurate and continuous forecasts.

The state vector $\mathbf{x} \in [0, 1]^N$ of the Kalman Filter contains an occlusion probability for each of the N forecasting time steps. In the prediction step of the filter, the previous forecast is advanced by the timespan between the acquisition of two consecutive images. This is done by using a state transition matrix Φ of the form


$$\Phi = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad (2)$$

that shifts the entries of the state vector \mathbf{x} . During a prediction step from one point in time to the next, the last entry will become the second last, and so on. The new last entry is initialized with zero occlusion probability.


In the update step, the predicted state vector is modified using a measurement $\mathbf{z} \in [0, 1]^N$. The measurement




(a) An example rectified sky image, with the position of the sun (red cross), and the part of the image that is used to establish the forecast (green central angle) marked. The dashed red line and the angle γ correspond to the axes in (b) and (c).



(b) Polar transform of the green circular sector.



(c) Cloud chart, i. e., segmented clouds in the polar transform.



(d) 10 minute forecast extracted from the cloud chart.

Figure 6: An overview over the forecasting method. From the input sky image (a) the circular region moving towards the sun is extracted and the polar transform of the sky image (b) and the segmented image (c) is calculated. The forecast (d) is calculated by taking a weighted sum along the γ -axis in (c).

vector consists of the preliminary forecast computed from the cloud chart. Since the preliminary forecast may be sparse, not all entries of \mathbf{z} need to be set. When this is the case, the values of the state vector \mathbf{x} are kept constant during the update step for missing measurement values. Furthermore, the covariance matrix for the measurements assigns higher measurement-error values to shorter forecasting times. More specifically, we use an error function that decreases exponentially with increasing forecasting time, i. e., distance to the sun.

The final forecast then becomes the current state of the Kalman filter after the update step. The advantages of the Kalman filter are that sparse forecasts can be combined to get a continuous one and that it is possible to extrapolate forecasts for short forecasting times. A direct forecast is often not possible for these short times because of saturated pixels close to the sun. In this case, previous forecasts are advanced to these shorter forecasting times using the state transition matrix in equation (2). Furthermore, the Kalman filter has the additional effect of smoothing the overall forecast. In Fig. 7 an example is shown on how the Kalman filter has combined several of the sparse preliminary forecasts to establish a continuous one.

Note that one of the preconditions for using a Kalman filter is that the underlying system model is linear. In our case this is satisfied, because the state transition models the evolution of time, which is linear. The possibly non-




Figure 7: An example on how the Kalman filter has combined several sparse preliminary forecasts to an optimal continuous one.

linear motion of the clouds is incorporated in the computation of the measurement vectors, when the distance of clouds to the sun is converted to arrival times. Another precondition when using a Kalman filter to estimate values from several noisy measurements over time, is that there are no fast changes of the estimated values that are not described by the transition matrix. In our application, fast changes of the occlusion probability, e. g., when multiple individual clouds are passing over the sun, are captured in the entries of a state or measurement vector. Over time, the Kalman filter we use only has to track gradual changes of the state vector and we, therefore, elude this problem.

For predicting occlusions or the expected global horizontal irradiance, the occlusion probability is compared to a threshold Θ_{occl} to get a binary forecast of occlusions.

4. Evaluation

The experimental evaluation is divided into two parts. First, we perform a sensitivity analysis to determine the optimal parameters for our forecasting method. This includes determining the cloud segmentation threshold Θ_{rb} , the opening angle of the circular region γ_{max} and the occlusion threshold Θ_{occl} . As performance metrics we used Precision, Recall and F_2 Score. In the second part of the experiments we evaluate the key aspect of our method, the use of a dense motion field, using the Root Mean Square Error (RMSE) of the predicted GHI and the forecast skill (see section 4.2).

4.1. Dataset

The dataset used for the evaluation consists of continuous data collected over 15 days in June 2013, March 2014 and April 2014. The meteorological conditions present in the dataset contain completely overcast days, sunny days and days with a frequently changing grade of cloudiness. **During the part of the evaluation in June 2013 convective clouds were prevalent. From March 2014 only sunny days are used, and the weather in April 2014 was dominated by the passing of cold fronts. From the days used in the evaluation, only on April 17 the weather was dominated by an anticyclone that was situated over Germany between the passing of two cold fronts.** Images and irradiance data was collected using the setup described in section 3.2. The camera and pyranometers were installed near Kitzingen, in Bavaria, Germany. The acquisition of the pyranometer irradiance data is synchronized with the image capture, i. e. data is recorded in 5 second intervals. To show the variability of the irradiance on these days, the pyranometer data for each day is shown in Appendix C. For our evaluation, night time data was discarded. All experiments are performed on the entire set of 15 days.

For quantitative evaluation we need to process the irradiance data to obtain timespans when the sun is occluded. Besides the measured pyranometer data we use an irradiance model that employs the position of the sensor and the time of day. The irradiance I is based on the model in Wong and Chow (2001) and is given by

$$I = I_{sc} \cdot (1 + 0.033 \cdot \cos(\Gamma)) \cdot 0.7^{m_a} \cdot \sin \alpha_s, \quad (3)$$

where $I_{sc} = 1367 \text{ W m}^{-2}$ is the solar constant, $\Gamma = 2\pi \left(\frac{N-1}{365}\right)$ is the day-angle of day N of the current year, and α_s is the solar altitude. The air mass m_a is given by

$$m_a = \left(\sin(\alpha_s) + 0.15(\alpha_s + 3.885)^{-1.253} \right)^{-1}. \quad (4)$$

The model is slightly simplified in comparison to Wong and Chow (2001) (higher order terms of Γ are neglected). However, it still matches the daily course of the irradiance properly.

To determine occlusions, we calculate the clear sky index, i. e. the fraction of the measured irradiance and the




Figure 8: Determining occluded (red) and unoccluded (blue) timespans from the forecasts using a preset occlusion threshold (---).

		Prediction	
		occluded	free sky
State	occluded	hits	misses
	free sky	false alarms	correct negatives

Table 1: Basic states used in computing Precision and Recall.

irradiance under clear sky conditions. The sun is assumed to be occluded when this ratio falls below a threshold of 0.65. We use this information as the ground truth data for the occlusion state in our evaluation.

4.2. Evaluation metrics

For evaluating the effects of the parameters in our method we use evaluation metrics that focus on the prediction of the state of the visibility of the sun. Precision, Recall and F_2 Score are used for measuring the accuracy of the prediction of this state. After determining the optimal parameters using these metrics, we also compare the effects of using a global cloud speed or the dense displacement field to establish the forecast. This is done using the Root Mean Squared Error (RMSE) between forecasted and measured GHI and the forecast skill as metrics.

4.2.1. Occlusions: Precision and Recall

The occluded timespans for the current forecast can be calculated by thresholding the forecast. The used threshold Θ_{occl} may either be a fixed value for all forecasting times, or a function of the forecasting time $\Theta_{occl}(t_{fc})$. The latter can be used to take into account that for larger forecasting times, a larger part of the sky is used to establish the forecast. Hence, a larger threshold is needed. In Fig. 8, we show an example for determining the occluded timespans from a forecast using a linear thresholding function.

Precision and Recall are metrics used in pattern recognition to evaluate classification tasks. They are applicable to our evaluation, since the prediction of the state of visibility of the sun can also be seen as a binary classification problem. Four combinations of predicted and actual state are possible (cf. Table 1), which are defined as the hits, correct negatives, misses and false alarms. For a large number of samples, Precision and Recall are defined as

$$PR = \frac{\text{hits}}{\text{hits} + \text{false alarms}} \quad (5)$$

$$RE = \frac{\text{hits}}{\text{hits} + \text{misses}}. \quad (6)$$

Precision describes how well predictions of occlusion matched real occlusions, while Recall describes how likely it is that real occlusions are actually predicted. In our application, Precision and Recall values can be calculated for each forecasting time individually, leading to Precision and Recall over forecasting time plots as shown in Fig. 9.

Although both Precision and Recall are easy to interpret in our case, in order to make decisions about the parameter choices it is preferable to use a single metric that combines both values. A common way of doing this is by using the F_β Score

$$F_\beta = (1 + \beta^2) \cdot \frac{PR \cdot RE}{\beta^2 \cdot PR + RE}, \quad (7)$$

which for $\beta = 1$ is the harmonic mean of Precision and Recall. We want to put more emphasis on the recall values and therefore choose to use the F_2 Score in our evaluation.

4.2.2. Root Mean Square Error (RMSE)

We calculate the RMSE between the predicted and measured irradiance. For calculating the predicted values we use the irradiance model shown in equation (3). We first determine the occluded timespans as described in the previous section. We then assume that the irradiance decreases to 35% of the current clear sky value I during an occlusion of the sun. The magnitude of this irradiance drop-off was determined by fitting the irradiance model to occluded time spans manually. The predicted irradiance is therefore calculated by multiplying the irradiance model by a factor of 0.35 or 1.0, depending on the state of occlusion. The RMSE can then be calculated for all forecasting times t_{fc} individually by

$$RMSE_{t_{fc}} = \sqrt{\frac{1}{T} \sum_t (I_{pyr}(t) - I_{t_{fc}}(t))^2} \quad (8)$$

where $I_{pyr}(t)$ is the measured irradiance and $I_{t_{fc}}(t)$ is the predicted irradiance for forecasting time t_{fc} , both at time t . The sum is over all T samples used in the evaluation.

4.2.3. Forecast skill

The forecast skill is used to compare our forecasts to a baseline method, namely persistence forecasting. The latter assumes that the current state of occlusion is not going to change for the forecasting time in question. We calculate the persistence forecast by first calculating the occlusion state from the pyranometer data (cf. section 4.1) and then shifting it by the forecasting time. Like for the other forecasts, the theoretical irradiance (for a clear sky) is then multiplied by the shifted occlusion state like in the previous subsection.

In order to compare the two forecast methods, we calculate the RMSE for our forecast ($RMSE_{t_{fc}}$) and the persistence forecasting method ($RMSE_p$). Following Quesada-Ruiz et al. (2014), the forecast skill $s(t_{fc})$ is then calculated




Figure 9: Precision and Recall for different values of the cloud segmentation parameter Θ_{rb} . The curves are calculated on the entire evaluation dataset.

as

$$s(t_{fc}) = 1 - \frac{RMSE_{t_{fc}}}{RMSE_p} \quad (9)$$

for each forecasting time t_{fc} . Skill values greater than zero imply that the proposed forecast outperforms the persistence method.

4.3. Sensitivity analysis

We first perform a sensitivity analysis to determine the optimal values of the parameters of our method.

4.3.1. Cloud segmentation threshold Θ_{rb}

The first parameter we evaluate is the threshold Θ_{rb} for cloud segmentation. As described in section 3.4, clouds are detected by thresholding the red/blue ratio of the pixels. Since white clouds are expected to have a value close to 1.0, we limit the range of the threshold parameter Θ_{rb} to values between 0.95 and 1.0. The resulting Precision and Recall plots for three values of $\Theta_{rb} \in \{0.95, 0.975, 1.0\}$ in this range are shown in Fig. 9.

All of the Precision plots are approximately constant over different forecasting times t_{fc} . A higher cloud threshold, however, results in a higher Precision value. The Recall plots show a different trend. The Recall values decrease with increasing forecasting times. The best Recall values are achieved by a threshold of $\Theta_{rb} = 0.975$. For the highest value of $\Theta_{rb} = 1.0$, however, the recall values have degraded which indicates that not all clouds were segmented correctly.

Since the overall best results according to the F_2 score are achieved with a value of $\Theta_{rb} = 0.975$, we use this setting for all subsequent experiments.




Figure 10: Precision and Recall for different values of the opening angle of the cone γ_{\max} . The curves are calculated using the entire evaluation dataset.

4.3.2. Cone opening angle γ_{\max}

The second parameter we evaluate is the opening angle γ_{\max} of the circular region used in the forecast. We have chosen to evaluate two values of $\gamma_{\max} \in \{30^\circ, 60^\circ\}$. As can be seen from the Precision and Recall curves in Fig. 10, the larger opening angle leads to better results for larger forecasting times. We therefore choose $\gamma_{\max} = 60^\circ$ for the remaining evaluation.

4.3.3. Occlusion threshold Θ_{occl}

The occlusion threshold Θ_{occl} is used for calculating the occluded timespans from the forecasts as illustrated in Fig. 8. The impact of this threshold is shown in Fig. 11, where Precision, Recall and F_2 score curves for constant and linearly increasing thresholds are shown. As expected, a lower threshold value leads to a better detection of occlusions, while also increasing the number of false detections (i. e., higher recall and lower precision values). According to the F_2 score, the best results are achieved using a fixed threshold value of $\Theta_{\text{occl}} = 0.25$.

4.4. Irradiance RMSE and forecast skill

After determining the parameters of our method, we use the RMSE between the automatic forecast and the measured GHI and the forecast skill for the final evaluation. **For this evaluation step the following parameters are used: $\Theta_{\text{rb}} = 0.975$, $\gamma_{\max} = 60^\circ$ and a fixed occlusion threshold $\Theta_{\text{occl}} = 0.25$.**

The main goal of this evaluation is to show the difference between using a global cloud speed and the dense motion vector field used by our method. For this experiment, the global cloud speed was assumed to be the mode of the distribution of motion vectors (cf. section 3.3). The

resulting RMSE and forecast skill values for the entire evaluation period are shown in Fig. 12. The RMSE using the dense motion field is lower than the RMSE for a global cloud motion for almost all forecasting times. This is also reflected in the forecast skill, where the skill value of our method exceeds the skill of the forecast using a global motion vector.

Using our proposed method that relies on non-rigid registration to determine a dense motion field, and a Kalman filter to track forecasts over time, we are able to establish forecasts of the GHI that outperform the persistence method even for short forecasting times of $t_{fc} = 60$ s. The RMSE our method achieves lies between 155 W m^{-2} and 200 W m^{-2} on the whole evaluation period.

These results indicate that the local deformations of clouds affect the reliability of irradiance forecasts. By considering local variations, the proposed method is more flexible in accounting for these deformations.


5. Conclusions

We have presented a novel method for forecasting occlusions of the sun, and thereby the GHI for solar power plants. Our method establishes continuous forecasts for occlusions up to a maximum forecasting time of 10 minutes. This distinguishes our work from previous forecasting systems, which only provided forecasts for a small number of fixed forecasting times.


A key element of our forecasting pipeline is the use of a Kalman filter, which improves the quality of forecasts by incorporating information from previous forecasts. The use of the Kalman filter also makes short term forecasts for forecasting times below 3 min possible. A direct forecast is difficult to calculate for these times because the region around the sun is usually overexposed or occluded by a shadow-band.

We further improved our forecasts by using a dense vectorfield of cloud displacement vectors that was calculated using a non-rigid registration algorithm. Instead of using only sparsely distributed motion vectors and a global cloud speed, this enables us to establish more accurate irradiance forecasts by determining the cloud speed for each distance to the sun individually.

For evaluating the proposed method we complemented the existing metrics of irradiance RMSE and forecasting skill by a sensitivity analysis based on Precision, Recall and the F_2 Score. The latter are metrics commonly used in pattern recognition for evaluating binary classification tasks and were used to find an optimal set of parameters for our method. We achieve a RMSE value for predicting the GHI on our dataset consisting of 15 days of challenging weather conditions between 155 W m^{-2} and 200 W m^{-2} , depending on the forecasting time. The forecast skill implies an improvement over the persistence baseline method for forecasting times longer than 60 s on the dataset used in our evaluation.



(a) Constant thresholds



(b) Linearly increasing thresholds


Figure 11: Precision, Recall and F_2 Score showing the influence of the occlusion threshold Θ_{occl} for constant and linearly increasing thresholds.

Figure 12: RMSE and forecast skill determined on the whole evaluation period showing the difference between using a global cloud speed and the dense motion vector field.

Appendix A. Demons non-rigid registration

The Demons [Thirion \(1998\)](#) algorithm is a widely-used methodology for performing non-rigid registration with a regularization based on diffusion.

For deriving the update equations of the Demons algorithm, we start by substituting $I_{t+\Delta t}(\mathbf{p} + \mathbf{d})$ by its first order Taylor expansion around $\mathbf{d} = (0, 0)^T$ in equation (1)

$$I_t(\mathbf{p}) - I_{t+\Delta t}(\mathbf{p}) = \nabla I_t(\mathbf{p}) \cdot \mathbf{d} \quad (\text{A.1})$$

with $\nabla I = (\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y})^T$ and $\mathbf{d} = (d_x, d_y)^T$. For registration purposes, the projection of the deformation field onto the image gradient is used

$$\mathbf{d} = \frac{(I_t(\mathbf{p}) - I_{t+\Delta t}(\mathbf{p})) \nabla I_t(\mathbf{p})}{\|\nabla I_t(\mathbf{p})\|^2}. \quad (\text{A.2})$$

As this equation becomes unstable for small values of $\|\nabla I_t(\mathbf{p})\|$, it is further normalized for the Demons algorithm

$$\mathbf{d} = \frac{(I_t(\mathbf{p}) - I_{t+\Delta t}(\mathbf{p})) \nabla I_t(\mathbf{p})}{\|\nabla I_t(\mathbf{p})\|^2 + (I_t(\mathbf{p}) - I_{t+\Delta t}(\mathbf{p}))^2}. \quad (\text{A.3})$$

In the Demons algorithm the deformation field is calculated iteratively using

$$\mathbf{d}^{n+1}(\mathbf{p}) = \mathbf{d}^n(\mathbf{p}) + \frac{(I_t(\mathbf{p}) - I_{t+\Delta t}(\mathbf{p} + \mathbf{d}^n)) \nabla I_t(\mathbf{p})}{\|\nabla I_t(\mathbf{p})\|^2 + (I_t(\mathbf{p}) - I_{t+\Delta t}(\mathbf{p} + \mathbf{d}^n))^2}. \quad (\text{A.4})$$

In principle, this equation could be solved for each pixel individually. However, this might lead to inconsistencies in the resulting deformation field. This is alleviated by blurring the deformation field \mathbf{d}^n with a Gaussian Kernel

after each iteration. As Thirion (1998) has shown, this corresponds to a regularization of the deformation field based on diffusion.

Since the Taylor expansion in equation (A.1) is only valid for small values of the deformation field, a multi-resolution approach has to be used when dealing with large deformations. In this approach, a pyramid of images with decreasing resolution is formed. Starting from the coarsest level (i. e. the smallest image), the deformation field is calculated. It is then resampled to match the size of the next higher level and used as the initialization for the calculations at that level. We use such a multi-resolution pyramid that starts with an image that is subsampled by a factor of 8.

Appendix B. Kalman filter

The Kalman filter is designed to produce optimal estimates of the state of a linear dynamic system, given inaccurate measurements. We will give a short introduction to its concepts. For a detailed discussion of the topic one may refer to the classic book of Maybeck (1979). The connection between measurements and the state is established by a measurement model. The dynamic system itself is described by a system model.

The system model describes the behavior of the state \mathbf{x}_t of the dynamic system over time. The Kalman filter assumes a linear model

$$\mathbf{x}_t = \Phi_{t-1}\mathbf{x}_{t-1} + \boldsymbol{\eta}_{t-1} \quad (\text{B.1})$$

where Φ_{t-1} is the transition matrix and $\boldsymbol{\eta}_{t-1}$ the additive system noise. The transition matrix may be time-dependent to model more complex systems, but in our application it is assumed to be constant. The measurement model is also described by a linear equation

$$\mathbf{z}_t = H_t\mathbf{x}_t + \boldsymbol{\nu}_t \quad (\text{B.2})$$

where the measurement matrix H_t connects the measurement \mathbf{z}_t to the state \mathbf{x}_t of the system, and $\boldsymbol{\nu}_t$ is an additive noise accounting for the uncertainty of the measurements. Again, the measurement matrix may be time-dependent, but is assumed constant in our application.

Given the two matrices which model the system's behavior, an optimal estimate of the state is calculated by alternating between a prediction and an update step. During the prediction step, the system model is used to predict the state \mathbf{x}'_t and its covariance P'_t at time t

$$\mathbf{x}'_t = \Phi\mathbf{x}_{t-1} \quad (\text{B.3})$$

$$P'_t = \Phi P_{t-1} \Phi^T + Q_{t-1} \quad (\text{B.4})$$

where Q_{t-1} is the covariance of the process noise $\boldsymbol{\eta}_t$ (assumed to be white, Gaussian noise). When the next measurement \mathbf{z}_t is taken, the update step is performed, which combines the measurement with the predicted value. In

this step, the measured value is weighted by the Kalman gain

$$K_t = \frac{P'_t H^T}{H P'_t H + R_t} \quad (\text{B.5})$$

where R_t is the covariance of the measurement noise $\boldsymbol{\nu}_t$ (which is again assumed to be white Gaussian noise). The estimate of the state and its covariance is then calculated

$$\mathbf{x}_t = \mathbf{x}'_t + K_t(\mathbf{z}_t - H\mathbf{x}'_t) \quad (\text{B.6})$$

$$P_t = (I - K_t)P'_t(I - K_t)^T + K_t R_t K_t^T \quad (\text{B.7})$$

We use a simpler version of the Kalman Filter, in which the measurement matrix is the identity matrix $H = I$, which slightly simplifies equations (B.2), (B.5) and (B.6).

Appendix C. Irradiation Data

The GHI for the evaluation period are shown in Fig. C.13 and Fig. C.14.

References

- Bernecker, D., Riess, C., Christlein, V., Angelopoulou, E., Hornegger, J., 2013. Representation Learning for Cloud Classification. *Pattern Recognition*, 395–404.
- Bosch, J., Kleissl, J., 2013. Cloud motion vectors from a network of ground sensors in a solar power plant. *Solar Energy* 95, 13–20.
- Chow, C.W., Urquhart, B., Lave, M., Dominguez, A., Kleissl, J., Shields, J., Washom, B., 2011. Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. *Solar Energy* 85, 2881–2893.
- Chu, Y., Pedro, H.T., Coimbra, C.F., 2013. Hybrid intra-hour DNI forecasts with sky image processing enhanced by stochastic learning. *Solar Energy* 98, 592–603.
- Inman, R.H., Pedro, H.T., Coimbra, C.F., 2013. Solar forecasting methods for renewable energy integration. *Progress in Energy and Combustion Science*.
- Kalman, R.E., 1960. A New Approach to Linear Filtering and Prediction Problems. *ASME - Journal of Basic Engineering* 82, 35–45.
- Kazantzidis, A., Tzoumanikas, P., 2012. Cloud detection and classification with the use of whole-sky ground-based images. *Atmospheric ...*
- Kidd, C., Levizzani, V., Bauer, P., 2009. A review of satellite meteorology and climatology at the start of the twenty-first century. *Progress in Physical Geography* 33, 474–489.
- Lonij, V.P., Brooks, A.E., Cronin, A.D., Leuthold, M., Koch, K., 2013. Intra-hour forecasts of solar power production using measurements from a network of irradiance sensors. *Solar Energy* 97, 58–66.
- Lorenz, E., Hurka, J., Heinemann, D., Beyer, H.G., 2009. Irradiance Forecasting for the Power Prediction of Grid-Connected Photovoltaic Systems. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 2, 2–10.
- Marquez, R., Coimbra, C.F., 2013. Intra-hour DNI forecasting based on cloud tracking image analysis. *Solar Energy* 91, 327–336.
- Maybeck, P.S., 1979. Stochastic models, estimation, and control. Number Bd. 1 in *Mathematics in Science and Engineering*, Academic Press, New York.
- Quesada-Ruiz, S., Chu, Y., Tovar-Pescador, J., Pedro, H., Coimbra, C., 2014. Cloud-tracking methodology for intra-hour DNI forecasting. *Solar Energy* 102, 267–275.
- Reikard, G., 2009. Predicting solar radiation at high resolutions: A comparison of time series forecasts. *Solar Energy* 83, 342–349.
- Thirion, J.P., 1998. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis* 2, 243–260.
- Wong, L., Chow, W., 2001. Solar radiation model. *Applied Energy* 69, 191–224.

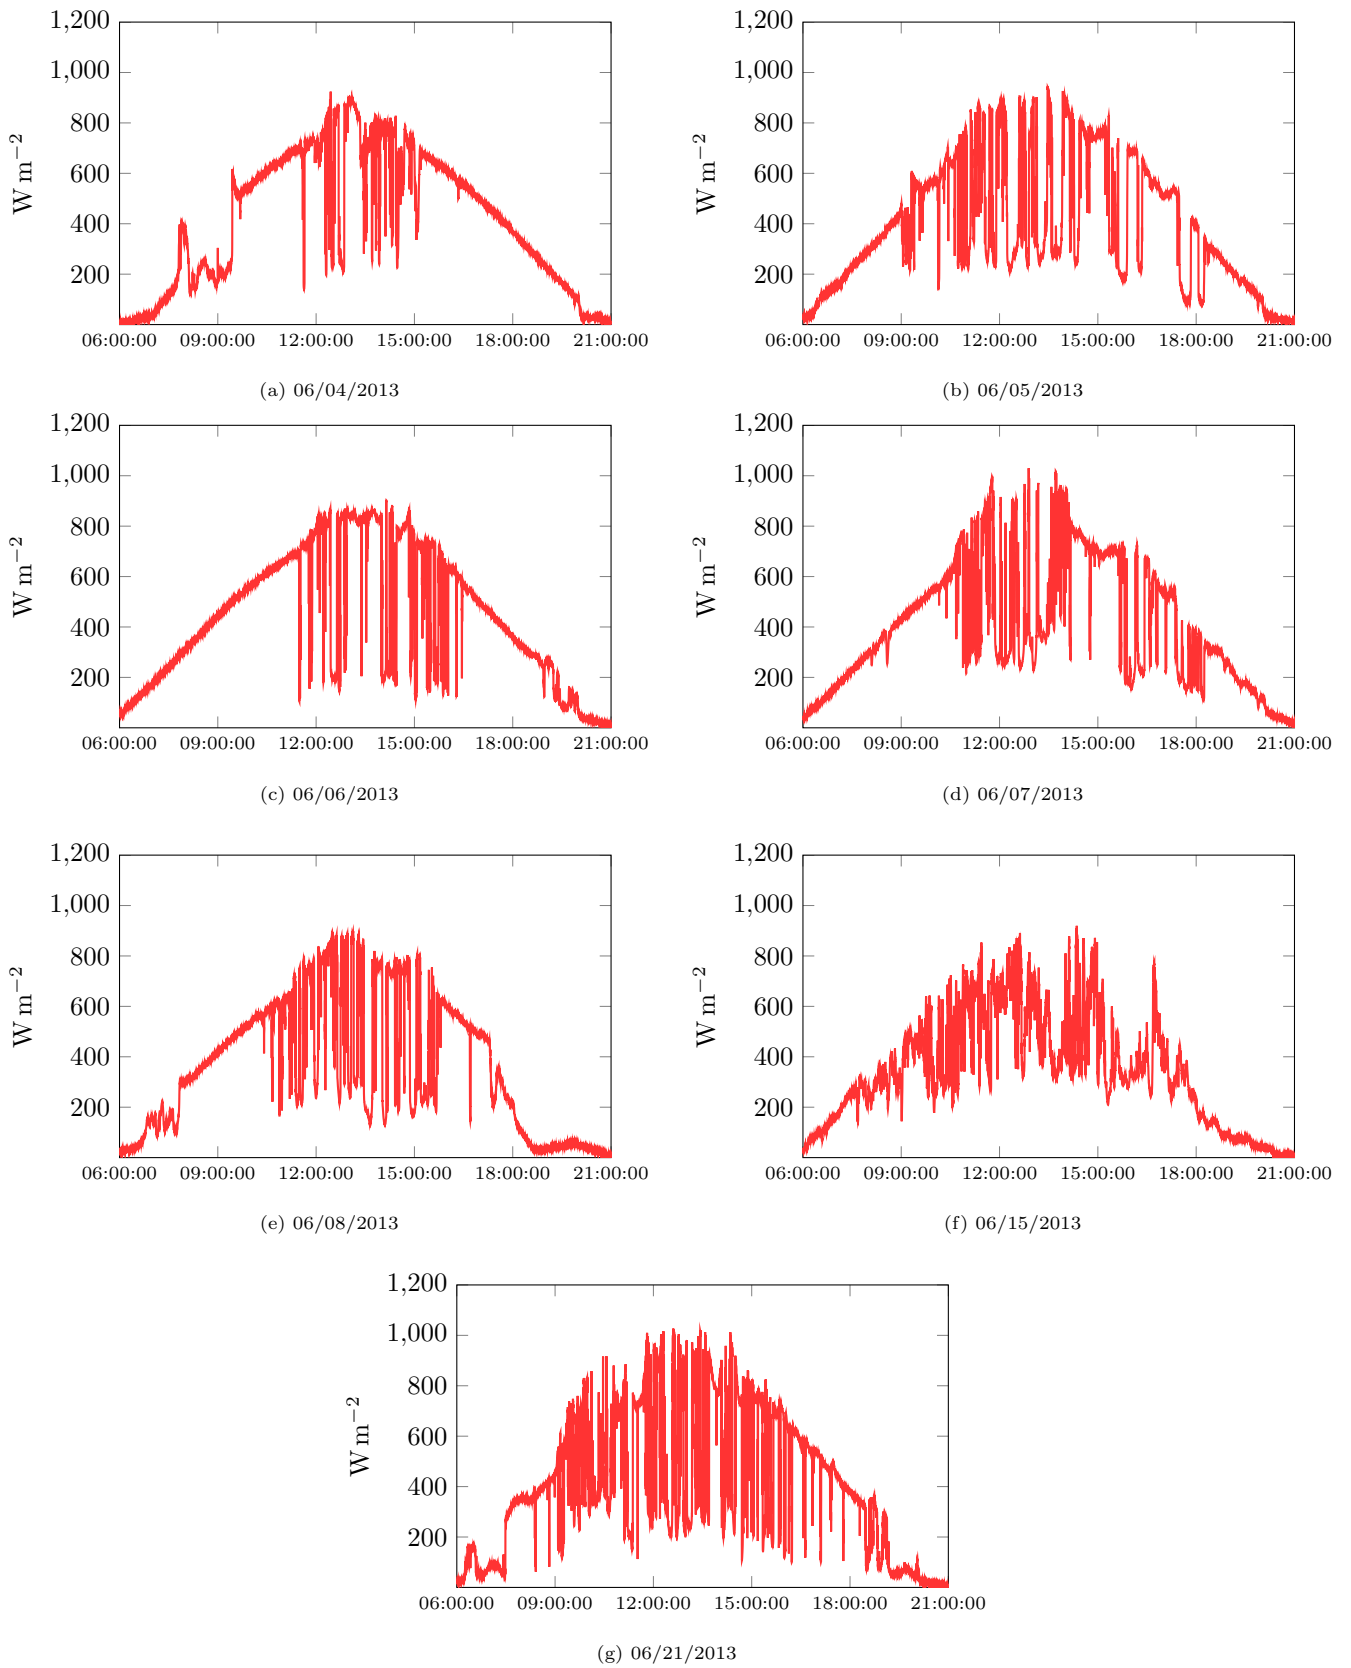


Figure C.13: Irradiation data from June 2013.

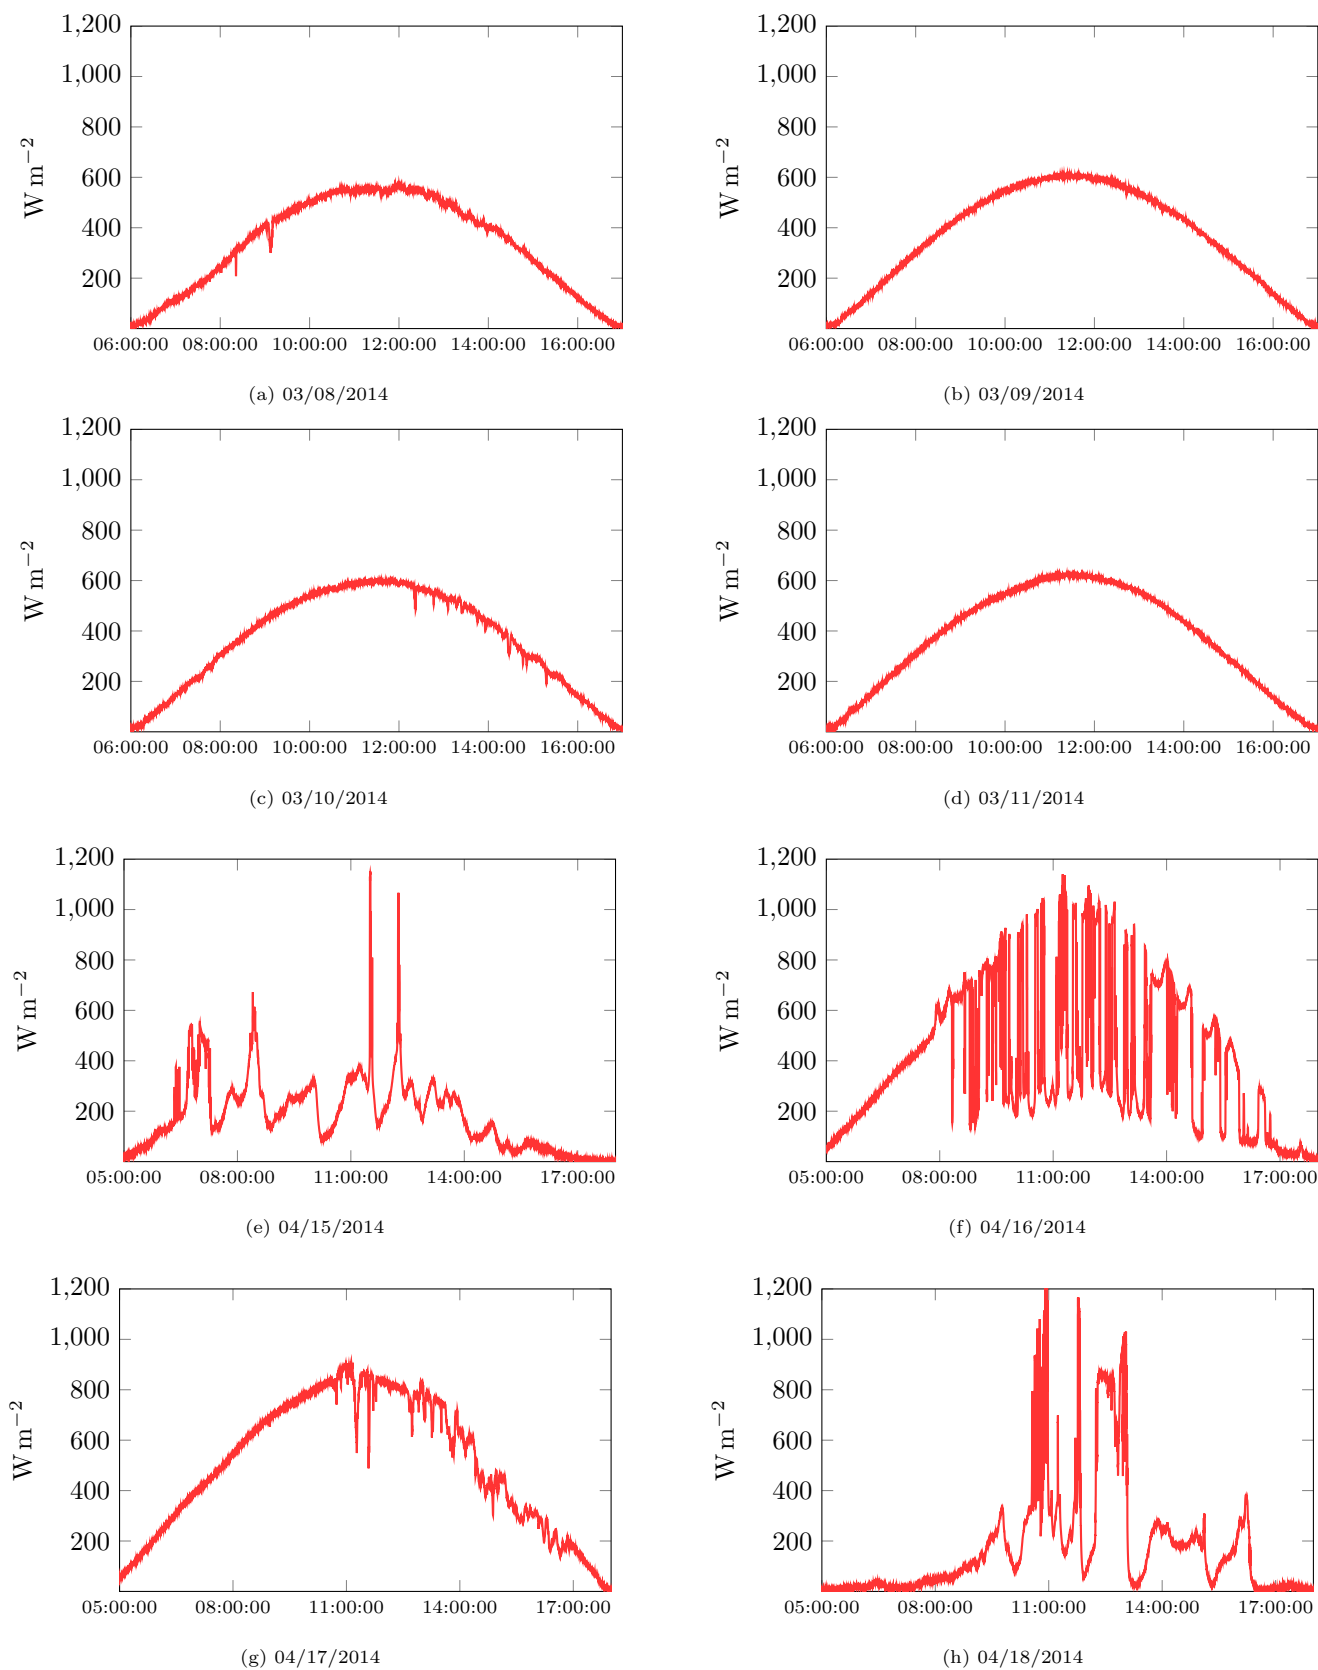


Figure C.14: Irradiation data from March and April 2014.