Real-time forecasting of solar irradiance ramps with smart image processing

Yinghao Chu, Hugo T.C. Pedro, Mengying Li, Carlos F.M. Coimbra [*]

Department of Mechanical and Aerospace Engineering, Jacobs School of Engineering, Center for Renewable Resource Integration and Center for Energy Research, University of California, 9500 Gilman Drive, La Jolla, CA 92093, USA

Abstract

We develop a standalone, real-time solar forecasting computational platform to predict one minute averaged solar irradiance ramps ten minutes in advance. This platform integrates cloud tracking techniques using a low-cost fisheye network camera and artificial neural network (ANN) algorithms, where the former is used to introduce exogenous inputs and the latter is used to predict solar irradiance ramps. We train and validate the forecasting methodology with measured irradiance and sky imaging data collected for a six-month period, and apply it operationally to forecast both global horizontal irradiance and direct normal irradiance at two separate locations characterized by different micro-climates (coastal and continental) in California. The performance of the operational forecasts is assessed in terms of common statistical metrics, and also in terms of three proposed ramp metrics, used to assess the quality of ramp predictions. Results show that the forecasting platform proposed in this work outperforms the reference persistence model for both locations.
© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Solar energy offers the opportunity to produce low- to zero-carbon power as solar generation technologies are now both functionally ready and nearly financially competitive to be deployed at large scales to the power grid. However, the variable nature of the solar resource remains an obstacle to achieving higher levels of grid penetration (Inman et al., 2013; Singh, 2013). At increasingly higher solar penetration levels, sudden power drops (chiefly caused by clouds) can adversely affect local power amplitude and frequency stability, which directly affects the quality of the power generated, and greatly affects the pricing of variable generation. One way to solve this issue, although expensive, is by storing electrical energy during sunny periods, which can be used during transient periods to smooth out power output. A low cost alternative that does not decrease the power output variability, as storage does, but decreases the uncertainty in the power output is solar forecasting.

With this goal in sight, different solar forecasting methods have been developed for various temporal horizons to meet the increasing demands for solar integration (Kalogirou, 2001; Li et al., 2008; Bacher et al., 2009; Huang et al., 2010; Mellit and Pavan, 2010; Hassanzadeh et al., 2010; Marquez and Coimbra, 2011; Pedro and Coimbra, 2012; Lave et al., 2012; Hart et al., 2012;

---

* Corresponding author.
E-mail address: ccoimbra@ucsd.edu (C.F.M. Coimbra).

Nomenclature

ANN	artificial neural network	MFR-7	multi-filter rotating shadowband radiometer
CI	cloud index	N	number
clr	clear-sky condition	NFV	number of forecast time points
CSL	clear-sky library	P	occurrence rate
CVM	cross validation method	p	persistence
DNI	direct normal irradiance	PIV	particle image velocimetry
FH	forecast horizon	RDI	ramp detection index
FRP	false ramp prediction	RM^*	non-dimensional ramp magnitude
FRI	false ramp prediction index	RMI	ramp magnitude index
FTM	fixed threshold method	RMSE	root mean square error
GA	genetic algorithm	SACI	smart adaptive cloud identification system
I	irradiance	s	forecast skill
ISFP	integrated solar forecasting platform	TNR	true no ramp prediction
k_e	excess kurtosis	t	time point
MBE	mean bias error	σ	standard deviation
MCE	minimum cross entropy method	μ_4	fourth central moment
MLP	multi-layer preceptron		

Marquez and Coimbra, 2013; Marquez et al., 2013a; Inman et al., 2013; Quesada-Ruiz et al., 2014; Chu et al., 2015). Among these models, artificial neural network (ANN) is one of the most widely used modeling and forecasting tools (Elizondo et al., 1994; Mellit and Kalogirou, 2008; Cao and Lin, 2008; Martín et al., 2010; Mellit and Pavan, 2010; Mellit et al., 2010; Marquez and Coimbra, 2011; Pedro and Coimbra, 2012; Marquez et al., 2013a; Marquez et al., 2013b; Chu et al., 2013). However, most of the cited studies do not consider cloud information as ANN inputs, while cloud information is one of the most important factors that cause ground-level irradiance ramps. Forecast models that consider cloud information are developed based on either remote sensing or local sensing. Remote sensing based models (Hammer et al., 1999; Perez et al., 2010; Marquez et al., 2013b; Nonnenmacher and Coimbra, 2014) that use satellite images have limited spatial and temporal resolutions and therefore not appropriate for intra-hour forecast (Ghonima et al., 2012; Marquez and Coimbra, 2013; Chu et al., 2013). To perform the intra-hour hour forecast, local sensing based models are developed using high resolution local sky images captured by high-frequency sky imagers (Chow et al., 2011; Marquez et al., 2013a; Marquez and Coimbra, 2013; Urquhart et al., 2013; Quesada-Ruiz et al., 2014). Typically, these sky imaging incorporated models do not employ stochastic learning techniques to statistically improve the robustness and accuracy of the forecasts. Forecast models that integrate ANN and cloud tracking methods have been developed in the recent years (Marquez et al., 2013a; Marquez et al., 2013b; Chu et al., 2013; Chu et al., 2014). For intra-hour horizons, these integrated methods achieve forecast skills ranging from 5% to 20% when tested on historical data.

With increasing local and global demand for solar energy generation and integration, the development of solar forecasting tools to mitigate power grid instability is now of critical importance, however, most of these forecast models are not assessed in real-time, focus on either global horizontal irradiance (GHI) or the direct normal irradiance (DNI), use expensive scientific or commercial sky imagers, and are evaluated only with statistical metrics that cannot explicitly assess forecast performance in predicting irradiance ramps.

In this work, we develop an operational Integrated Solar Forecasting Platform (ISFP) that takes cloud information as exogenous inputs to ANN and simultaneously forecasts one minute averaged values of both GHI and DNI ten minutes in advance. We use a generic off-the-shelf, high-resolution fisheye dome network camera to capture sky images for the ISFP. Compared to sky imagers used in other works (e.g. YES TSI-440/880), the network cameras have advantages such as substantially lower cost, higher spatial resolution, portability, and ease of installation.

Accurate forecast of irradiance ramps is essential for inverter control, plant management and real-time dispatch operations for solar power plants (Zhang et al., 2013; Florita et al., 2013). However, common statistics metrics, such as root mean square error (RMSE), are unable to quantitatively evaluate the performance of a forecast model in estimating irradiance ramps. Zhang et al. (2013, 2015) evaluate a suite of metrics which are commonly used to assess solar forecasts and suggest that skewness, kurtosis, and Renyi entropy are qualitatively sensitive to the performance of ramp forecasts. Florita et al. (2013) propose a swinging door algorithm to extract solar ramps by identifying the start and end points of ramps. However, these available metrics do not explicitly quantify the

performance of ramp forecasts. Therefore, to quantitatively assess the performance of ramp forecast considering both ramp duration and ramp magnitude, we propose three new metrics to assess the quality of ramp forecasts: the ramp detection index (RDI), the false ramp prediction index (FRI), and the ramp magnitude index (RMI). The definition of these three metrics is presented in Section 3.

The data used for the ISFP model development and operational forecasts is presented in Section 2. Detailed methods are presented in Section 3, including cloud detection, image processing to extract numerical cloud information, the training and optimization of the ANNs, and the metrics to assess the forecasting performance. Results and discussion are presented in Section 4, and conclusions are presented in Section 5.

2. Data

We collected irradiance data using 2 Multi-Filter Rotating Shadowband Radiometers (MFR-7) installed at our observatories in Folsom (latitude = 38.64°, longitude = −121.14°) and San Diego (latitude = 32.88°, longitude = −117.24°), California. Manufactured by Yankee Environmental Systems, Inc, the MFR-7 is able to simultaneously measure the GHI and DNI components of broadband solar irradiance every 30 s. Campbell Scientific CR1000 data loggers are used to log one minute average GHI and DNI values.

Sky images are obtained from 2 fisheye cameras installed next to the MFR-7s. The fish-eye cameras use a 3.1MP CMOS sensor and a 360° panoramic view lens. The cameras capture sky images and transfer them (via FTP) to a remote server every minute. Irradiance data and sky images are accessed in real-time and stored in a MySQL database.

High quality data sets are essential to generate high fidelity forecasts. The MFR-7s are first-class radiometers that meet the accuracy requirements of this study. Irradiance measurements are manually inspected and compared with measurements from Licor LI-200 pyranometers which are deployed close to the MFR-7s for data quality control. We clean the camera domes regularly to maintain the quality of sky images, and discard sky images that show significant amount of dust on the camera dome.

In addition, we only consider data points collected during periods for which the solar elevation angle is higher than 20°. There are three reasons for discarding periods of low solar elevation: (1) grid elements in the sky image processing (Section 3.2) may be placed outside the projected image and unable to provide the necessary inputs for solar forecasts; (2) ground obstacles such as trees and buildings (see Fig. 1) occasionally obscure the sun and decrease the accuracy of irradiance measurement and; (3) higher scattering of direct beam generates image glare (Chu et al., 2014), which adversely affects the accuracy of cloud detection and solar forecasts.

Six months (January 13, 2013 to July 10, 2013) of historical irradiance and sky-imaging data were collected in Folsom for the training of the stochastic component of the ISFP. The trained algorithm was used to generate operational real-time forecasts for both Folsom and San Diego by ingesting live data streams of irradiance and sky images into the model. The learning data include the possible range of irradiance variability. After the supervised learning process (described in Section 3), the ISFP model was deployed to forecast real-time GHI and DNI in Folsom every 10 min as an unattended task. To study the generality of the ISFP model, we duplicated and deployed as a robust operational forecast of GHI and DNI for the San Diego observatory.

The limited field of view of the sky camera prevents using this methodology for forecast horizons longer than 20–30 min depending on the cloud speed. Therefore in this work, the real-time forecast produces 10-min forecasts of GHI and DNI. In a previous work (Chu et al. (2014)), we have explored historical forecast with different forecast horizons (e.g. 5- and 15-min) for a similar methodology.

In order to determine the performance of the operational forecasts, we monitored and analyzed its output for the period of July 31, 2013 to February 2, 2014 in Folsom and from November 23, 2013 to February 11, 2014 in San Diego. For both locations, periods of real-time forecasts include the possible range of irradiance variability caused by diverse weather conditions. The statistical information (variability, occurrence and magnitude of ramps, etc.) for the irradiance time series is listed in Table 1. The variability mentioned in Table 1 is the standard deviation of the clear-sky index, which is defined as the ratio of measure irradiance to the clear-sky irradiance predicted by the clear-sky model (discussed in Section 3.4). San Diego region has a coastal micro-climate. The presence of marine layer, usually in the morning and late afternoon, increases the occurrence rate of solar ramps with different magnitudes. Therefore, the variability of irradiance in San Diego is statistically higher than that in Folsom.

3. Methods

In this section, we present several tools used to develop the ISFP and evaluate its performance. First, the cloud identification method (Section 3.1). Second, the image processing method to generate the cloud index (CI) time-series, which are the inputs to the ISFP (Section 3.2). Third, the ANN based ISFP and the genetic algorithm (GA) optimization of ISFP (Section 3.3). Forth, the performance assessment metrics for ISFP (Section 3.4).

3.1. Smart adaptive cloud identification

We develop a smart adaptive cloud identification system (SACI) (Chu et al., 2014) to conduct cloud detection for images captured by the fish-eye cameras used in this experiment. Firstly, each sky image is categorized as clear or
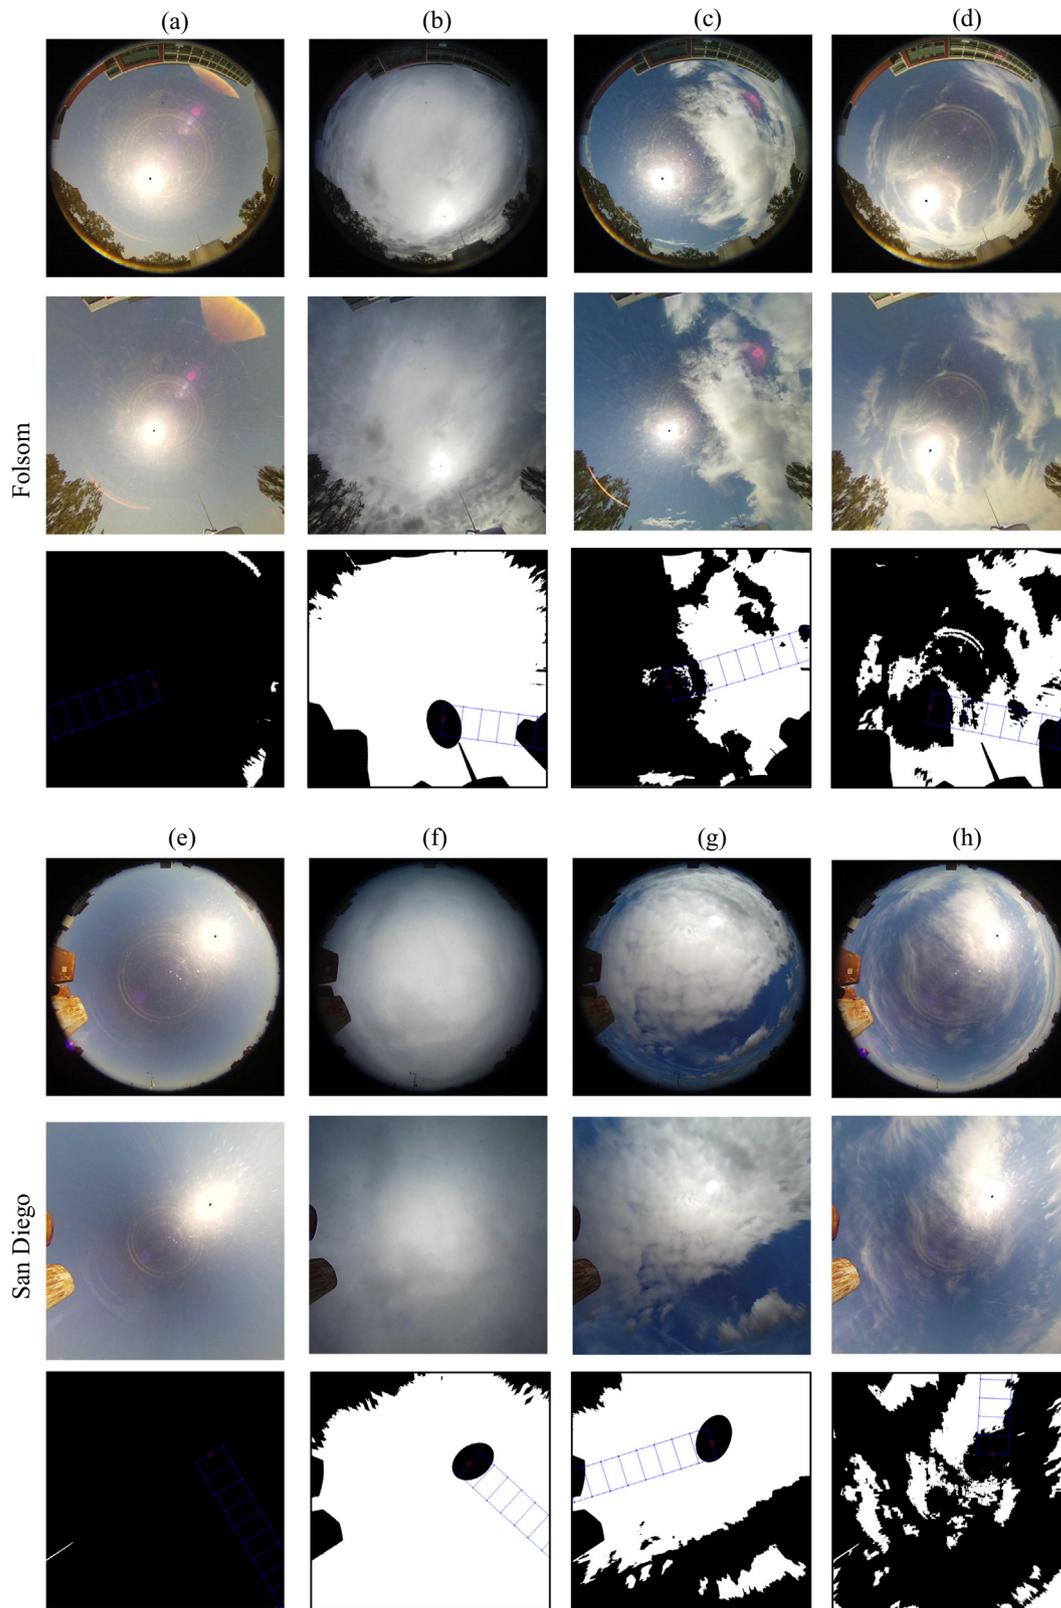
Fig. 1. Examples of SACI cloud identification. Sky images for Folsom: (a) clear period (2013-08-05), (b) overcast period (2013-11-19), (c) partly cloudy period with optically thick clouds (2013-10-06), and (d) partly cloudy period with optically thin clouds (2013-11-11). Sky images for San Diego: (e) clear period (2013-12-30), (f) overcast period (2014-01-30), (g) partly cloudy period with optically thick clouds (2014-02-09), and (h) partly cloudy period optically thin clouds (2013-12-01). The first and the fourth rows present the original images, the second and the fifth rows present the projected images, and the third and sixth rows present the detected binary cloud maps with the grid elements.

Table 1
Statistics of irradiance data collected during the analysis periods. Variability is the standard deviation of the 10-min changes of the clear-sky index; NFV is the number of forecast time points; RM* is the non-dimensional ramp magnitude (introduced in Section 3.4). P(RM*) represents the occurrence rate of ramp with specific non-dimensional magnitude.

Location	Irradiance	NFV	Variability	$P(0.1 < RM^* < 0.2)$ (%)	$P(0.2 < RM^* < 0.3)$ (%)	$P(0.3 < RM^* < 0.5)$ (%)	$P(RM^* > 0.5)$ (%)
Folsom	GHI	6581	0.1053	5.00	2.50	2.10	1.00
	DNI	6581	0.1197	3.00	2.10	1.80	1.70
San Diego	GHI	2483	0.1562	11.40	4.80	4.60	2.40
	DNI	2483	0.1683	8.50	3.90	4.70	3.10

cloudy using five criteria computed from past 10-min GHI time-series. The five criteria are mean GHI, max GHI, length of GHI time-series, variance of GHI changes, and maximum deviation from clear-sky gradient (see Long and Ackerman, 2000; Younes and Muneer, 2007; Reno et al., 2012 for more details). Secondly, if the image is categorized as cloudy, the hybrid thresholding method (Li et al., 2011) is used to analyze red blue ratio histograms and further categorize the image as either overcast or partly cloudy. Thirdly, after the image categorization, the SACI employs Fixed threshold method (Li et al., 2011) for overcast images, clear sky library method with fixed threshold (Ghonima et al., 2012) for clear images, and clear sky library method with adaptive threshold (Otsu, 1979; Li and Lee, 1993; Li and Tam, 1998) for partly cloudy images. Examples of SACI cloud detection are shown in Fig. 1.

3.2. Grid cloud fraction method

Clouds that move to shade the sun will significantly reduce DNI and therefore are more relevant to solar forecast. Given that ISFP requires numerical inputs, we apply a grid-cloud-fraction method to sky images (Marquez and Coimbra, 2013) to obtain numerical information regarding clouds that move towards the sun.

The grid-cloud-fraction method can be summarized in five steps (see Marquez and Coimbra, 2013 and Chu et al., 2013 for more details of this method):

1. Round sky images from fish-eye cameras are projected onto a flat rectangular grid to remove the geometric distortion of the images (see Fig. 1, second and forth row).
2. Pairs of consecutive images are processed using Particle Image Velocimetry (PIV) to compute the moving direction of the clouds (Mori and Chang, 2003). The PIV method provides average cloud velocity in the sky. However, PIV is unable to identify multiple layer clouds that move with different velocity.
3. SACI is used to identify clouds from sky images to produce binary cloud maps. Our current algorithm provides only binary cloud maps and therefore is unable to determine the transmittance of clouds.
4. A set of grid elements is placed in the reverse direction of the cloud movement from the sun position on the binary cloud map (see Fig. 1, third and sixth row). In this work, the dimensions of the grid elements is empirically set to 120 (length) × 200 (width) pixels.

5. Cloud indices CI_i are computed as the fraction of pixels identified as cloud in grid element i. The resulting CI_i time-series can then be used as input to the ANN-based ISFP.

3.3. Artificial neural network and genetic algorithm

Artificial neural network (ANN) is a classification and regression tool widely used in solar modeling and forecasting (Bishop, 1994; Mellit and Kalogirou, 2008; Marquez and Coimbra, 2011; Inman et al., 2013; Chu et al., 2013). In this work, we employ multilayer preceptron neural networks (MLP). Signal processing elements (neurons) are the processing units of the neural network. Neurons are placed in layers, and the layers between the input layer and the output layer are called hidden layers (Inman et al., 2013). On each hidden layer, the i-th neuron sum the weighted outputs X_j of the previous layer and add a bias or threshold β_{ij} to the sum. Then the sums are processed by the activation function of neurons $f(\cdot)$ (this work uses sigmoidal functions) to generate outputs Y_i, which work as the inputs for neurons on the following layer:

$$Y_i = f\left(\sum_{j=1}^{N}(w_{ij}X_j + \beta_{ij})\right), \tag{1}$$

where N is the total number of outputs received from previous layer, X_j is j-th output from previous layer, Y_i is the output of the i-th neuron on current layer, w_{ij} and β_{ij} are the weight and bias of the j-th input on the i-th neuron, and $f(\cdot)$ is the activation function. The MLP parameters (the weight w_{ij} and bias β_{ij}) are determined from learning data in a supervised training process. In this work, the training process is the Bayesian regularization process with Levenberg–Marquardt optimization (Bishop, 1994; Inman et al., 2013) using clear-sky indexes 10 min afterward as the training targets. To implement this method for other forecast horizon (e.g. 5-, or 15-min), the MLP needs to be retrained with updated targets. Once the training process is converged and validated, the MLP model is able to make prediction with new inputs.

The popularly used cross validation method (CVM) is employed to prevent over-fitting of the MLP (Kohavi, 1995; Jain et al., 2000; Geisser, 1993; Lendasse et al., 2003; Efron and Gong, 1983; Chu et al., 2013). Implementation of CVM includes the following three steps:

1. The training data is divided into N subsets. We set $N = 10$ as suggested in literature (Kohavi, 1995; McLachlan et al., 2004).
2. We train the MLP with randomly selected N-1 subsets and use the remaining one subset to assess the performance of the trained MLP model. We use root mean square error (RMSE) of MLP model predictions and targets as a criterion of MLP model performance.
3. Repeat the above two steps N times such that each subset is used as the validation set. We use the average of N RMSE as the performance criterion of MLP model. An over-fitted model is prone to have high average RMSE and optimized model should have the least average RMSE.

We employ the GA to identify the optimal MLP scheme (number of layers, number of neutrons on each layers, potential inputs) that achieves the smallest validation RMSE. GA is a tool to find the optimal solution from massive potential solutions (Holland, 1992; Mellit and Kalogirou, 2008; Pai and Hong, 2005; Marquez and Coimbra, 2011; Pedro and Coimbra, 2012; Chu et al., 2013). The GA optimization is initialized with randomly generated individuals. The GA objective function assesses each individual and returns a fitness value. Then the individuals with best fitness are selected as parents and generate the new generation through cross-over and mutation. The iterative process ends when the average fitness of the individuals converges. The average RMSE from CVM method is used as the GA objective function in this work.

In total, 12 variables are chosen as GA inputs. The first 10 values are binary numbers used to determine whether a possible input should be selected as inputs to the MLP. The first five of these inputs are the clear-sky indexes of measured irradiance values at the current time, 5 min ago, 10 min ago, 15 min ago and 20 min ago. We use clear-sky indexes as both inputs and targets to exclude the effect of the diurnal solar variation from the process of training and forecasting. The 6th input is the total sky cloudiness. The last four inputs are cloud indices $CI_1, \cdots, CI_4$, calculated by the grid fraction method (Section 3.2). The goal of GA is to identify which set of inputs are more useful to the forecasts. GA selected inputs (both lagged irradiance and CIs) are ingested by the MLP with fixed order. The 11th variable in the GA search-space is an integer that specifies the number of hidden layers for the MLP, and the 12th GA variable, also an integer, is the number of neurons per hidden layer in the MLP.

The GA optimization are run as parallel jobs using 10 cores (2.39 GHz processor, 24 GB RAM). Once we have the optimized MLP scheme via GA (indicated by the convergence of average GA fitness in about 50 generations), we train the GA-optimized MLP as stated at the beginning of this section and apply the trained MLP to real-time forecast. The optimization and training work takes about 24 h to complete.

3.4. Testing process

The persistence model is selected as a reference model to evaluate the performance of ISFP. The persistence model is the simplest forecast model and achieves high accuracy in low variability periods. For this model we assume that the clear-sky index remains constant within the forecasting interval $[t, t + FH]$:

$$\widehat{I}_p(t + FH) = \frac{I(t)}{I_{clr}(t)} \times I_{clr}(t + FH) \tag{2}$$

where $\widehat{I}_p(\cdot)$ is the prediction of persistence model, FH is the forecast horizon, $I(\cdot)$ is the measured GHI or DNI values, and $I_{clr}(\cdot)$ is the predicted GHI or DNI from a clear-sky model. Assuming the persistence of the clear-sky index is better than assuming persistence of irradiance because it removes the effect of the diurnal solar variation. In this work, we use the clear-sky model proposed by Ineichen and Perez (2002). The extraterrestrial irradiance value I_0, current time, longitude, latitude and average Linke Turbidity are the inputs for this clear-sky model. The detailed method to calculate Linke Turbidity is discussed in Ineichen and Perez (2002). The employed clear-sky model is less accurate in morning and afternoon, particularly for direct normal irradiance (Chu et al., 2013; Quesada-Ruiz et al., 2014). The error of the clear-sky model is visible in Figs. 3a and 4a.

We employ four popularly used statistical metrics to assess the performance of the ISFP. The first one is the mean biased error (MBE):

$$MBE = \frac{1}{m} \sum_{t=1}^{m} (\widehat{I}(t) - I(t)). \tag{3}$$

The second one is the root mean square error (RMSE):

$$RMSE = \sqrt{\frac{1}{m} \sum_{t=1}^{m} (\widehat{I}(t) - I(t))^2}. \tag{4}$$

The third one is the forecasting skill (s), which measures the improvement of the proposed forecast method (ISFP) over the reference persistence model:

$$s = 1 - \frac{RMSE}{RMSE_p}. \tag{5}$$

The forth one is the excess kurtosis, which quantifies the 'peakness' of a distribution and the 'heaviness' of its tail Kenny (1961), Joanes and Gill (1998), Chu et al. (2013),

$$k_e = \frac{\mu_4}{\sigma^4} - 3, \tag{6}$$

where μ_4 is the fourth central moment of the forecast errors, and σ is the standard deviation of the forecast errors.

Given that one of the main goals in this work is to assess the forecasting performance when ramps in the irradiance are present, we define a set of three metrics for this purpose. In the first place, we need to establish the definition

of a ramp and the criteria to identify a ramp. There are two common definitions of ramps: (1) the irradiance difference between the start and end points of a time interval, and (2) the difference between minimum and maximum irradiance within a time interval (Zheng and Kusiak, 2009; Kamath, 2010; Florita et al., 2013; Nonnenmacher et al., 2014). In this case, the two definitions are equivalent since in the interval $[t, t + FH]$ there are only two available data points: the measured valued at time t and the measured at time $t + FH$ for the measured ramps, and the measured valued at time t and the forecasted value at time $t + FH$ for the forecasted ramps. Ramps are identified when the absolute difference of irradiance between the start and end points of a time interval is higher than a threshold. The time interval is set equal to the 10 min forecast horizon.

Low-amplitude fluctuations in the measured irradiance time series may be caused by measurement errors, diurnal solar circle, and other unexpected factors, even during clear periods. Therefore, to analyze the significant ramps in irradiance that impair the stability of solar production, we empirically use a threshold equal to 10% of the current clear-sky irradiance I_{clr} to identify 'true' ramps. We use I_{clr} to determine the value of threshold for two reasons: (1) during periods of low solar elevation ($\leqslant 30°$), the average irradiance and corresponding ramp magnitude are relatively lower than their counterparts for periods of high solar elevation. Therefore, an I_{clr} based threshold is less likely to miss ramps compared to a fixed threshold during the low sun elevation period; (2) I_{clr} will not vary drastically during cloudy periods compared to a threshold based on measured irradiance.

The ramp analysis is always applied after the operation of ISFP. Therefore, the magnitude of the threshold has no effect on the predictions of ISFP. However, if a large threshold is used, solar ramps with relatively small magnitude may not be identified. On the other hand, if a small threshold is used, measurement errors may be erroneously classified as ramps. Therefore, when applied elsewhere, the magnitude of the ramp threshold may be adjusted depending on the uncertainty of measurements and maximum tolerance of ramp magnitude for the specific application.

The assessment of the ramp prediction is based in three values: the measured irradiance values $I(t), I(t + FH)$, and the prediction $\widehat{I}(t + FH)$ made at t for time $t + FH$. We define the value of a ramp as the difference between the two measured irradiance ($I(t + FH) - I(t)$), and define the ramp prediction as the difference between the predicted irradiance and the measured irradiance ($\widehat{I}(t + FH) - I(t)$). An 'up ramp' has value greater than $10\% I_{clr}$ and a 'down ramp' has value smaller than $-10\% I_{clr}$. Otherwise, we define there is no ramp. Six possible outputs of the ramp forecasts are presented in Fig. 2. The ramp magnitude (RM, W/m^2) is defined as the absolute of ramp value ($|I(t + FH) - I(t)|$). For simplicity of following analysis, we define a non-dimensional ramp magnitude as RM$^* =$ RM/I_{clr}.

We propose three metrics to assess the quality of the ramp forecasts. The first one is ramp detection index (RDI). When a ramp occurs (the ramp magnitude exceeds the $10\% I_{clr}$ threshold), the ramp prediction is denoted as a 'hit' if (1) the magnitude of the ramp prediction exceeds the ramp threshold (RM$^* > 0.1$) and (2) ramp prediction and
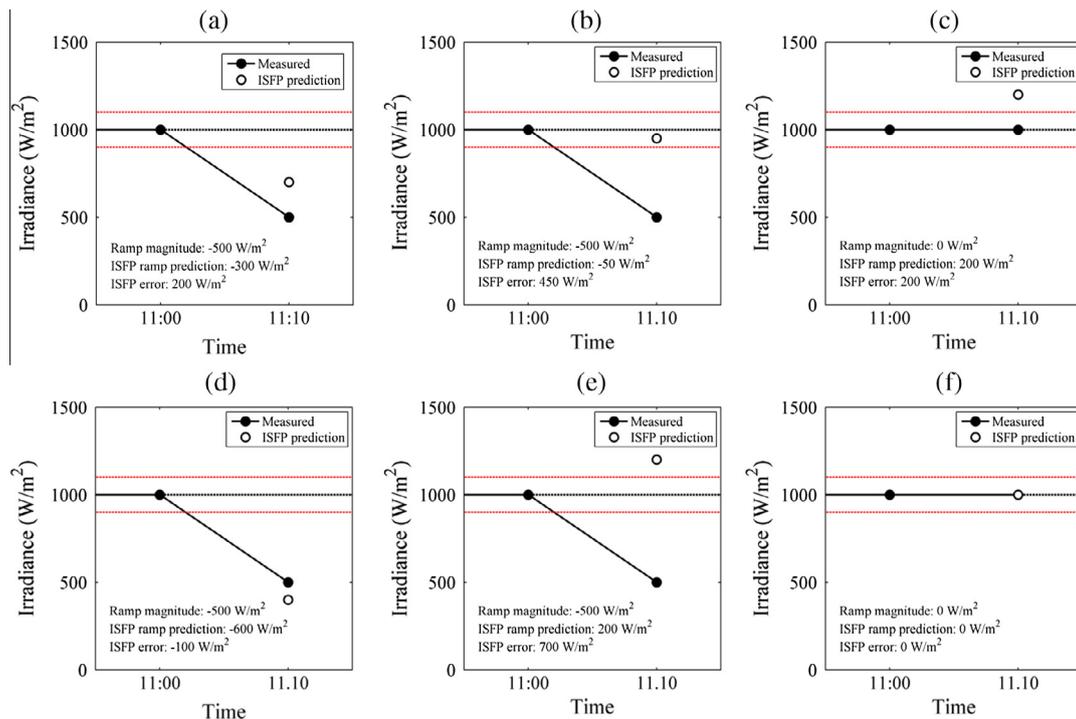

Fig. 2. The six possible outputs for ramp forecast when $I_{clr} = 1000$ W/m^2.

the value of ramp have the same sign $\frac{I(t+FH)-I(t)}{\widehat{I}(t+FH)-I(t)} > 0$. Examples of 'hits' are illustrated in Fig. 2(a and d). Otherwise, the ramp prediction is denoted as a 'miss' as in Fig. 2(b and e). The RDI is defined as the occurrence rate of ramp 'hits':

$$\text{RDI} = \frac{N_{hit}}{N_{hit} + N_{miss}}. \tag{7}$$

The second one is false ramp index (FRI). False ramps occur if, for time intervals where no ramp was observed, the predicted $\widehat{I}(t + FH)$ is beyond the range of $I(t) \pm 10\% \, I_{clr}$ (see Fig. 2c). Otherwise, it is a true no ramp prediction (TNR). Such events are usually observed during clear or overcast periods (see Fig. 2f). Once these values are quantified the FRI is defined as the occurrence rate of the false ramp predictions:

$$\text{FRI} = \frac{N_{FRP}}{N_{FRP} + N_{TNR}}. \tag{8}$$

The last metric proposed is the ramp magnitude forecast index (RMI). Using RDI and FRI, we could quantitatively assess the forecast accuracy in predicting the occurrence of ramps. However, these two metrics do not quantify the performance of forecast platform in predicting the magnitude of ramps. RMI is defined as the skill of the forecast in predicting the magnitude of ramps:

$$\text{RMI} = 1 - \sqrt{\frac{\sum_{i=1}^{N_r}(I(t_i + FH) - \widehat{I}(t_i + FH))^2}{\sum_{i=1}^{N_r}(I(t_i + FH) - I(t_i))^2}} \tag{9}$$

where N_r is the number of ramp events, $I(\cdot)$ is the measured value and $\widehat{I}(\cdot)$ is the predicted value. Higher RMI values indicates a more accurate prediction of ramp magnitude, with RMI = 1 representing a perfect prediction of ramp magnitude.

4. Results and discussion

4.1. GA optimization

The results of GA optimization are presented in Table 2, which includes the optimal set of input variables, number of hidden layers, and the number of neurons per layer. Table 2 shows that the total sky cloud coverage and the cloud indices are selected as inputs to both DNI and GHI forecasts. Therefore, cloud cover information is useful

for short-term solar forecasts. These optimization results match our earlier studies based on historical data in Chu et al. (2013) and Chu et al. (2014), which suggest that cloud information increases the forecast skills of MLP-based models by 3–8% depending on forecast horizon and weather conditions.

4.2. Error metrics

Results for the operational forecasts in real-time are presented in Table 3. Both persistence and ISFP models show small MBE (less than 5 w/m^2). In terms of RMSE and forecasting skill s, ISFP outperforms the persistence model. The higher forecasting skill is attributed to the combination of cloud tracking and stochastic learning. Improvements over the persistence model are in the range of 6.0–11.3% depending on location and irradiance component. The ISFP results show smaller excess kurtosis k_e than persistence model and the results are in agreement with the discussions in Chu et al. (2013): high kurtosis distributions have sharper peaks and longer, heavier tails. The MLP-based ISFP produces fewer instances of large errors and more instances of moderate errors than persistence model, resulting in an error distribution with shorter tail and rounder shoulders,resulting in a smaller kurtosis.

Example time-series of forecast and corresponding error are shown in Figs. 3 and 4 for Folsom and San Diego, respectively. Because the ISFP operates to provide one prediction every 10 min, each ISFP prediction is used to represent the expected level of irradiance for a 10-min interval in these two figures. Each row of these two figures represents one of the typical weather conditions: clear, overcast, partly cloudy with optically thin clouds, and partly cloudy with optically thick clouds. Representative sky images corresponding to Fig. 3a–d are shown in Fig. 1(a–d). Representative sky images corresponding to Fig. 4 are shown in Fig. 1(e–h).

During low irradiance variability periods (clear and overcast), both persistence and ISFP models achieve excellent performance with relatively low errors. During cloudy periods, ISFP achieves positive forecast skills over the persistence forecasts in both locations. The possible error sources of ISFP include but are not limited to: (1) the cloud

Table 2
GA optimized variables for the GHI and DNI forecasts. Indices represent selected MLP inputs: 1–5 represent the time-lagged irradiance values, 6 represents the total sky cloudiness, and 7–10 represent the cloud indices CI_i. L represents the number of hidden layers, and N represents the number of neurons per hidden layer.

Irradiance	Indices of inputs	L	N
GHI	1 2 3 5 6 7 8 9 10	1	8
DNI	1 2 3 4 6 7 8 9 10	1	10

Table 3
Statistical metrics for real-time GHI and DNI forecasts in Folsom and San Diego for 10 min horizon.

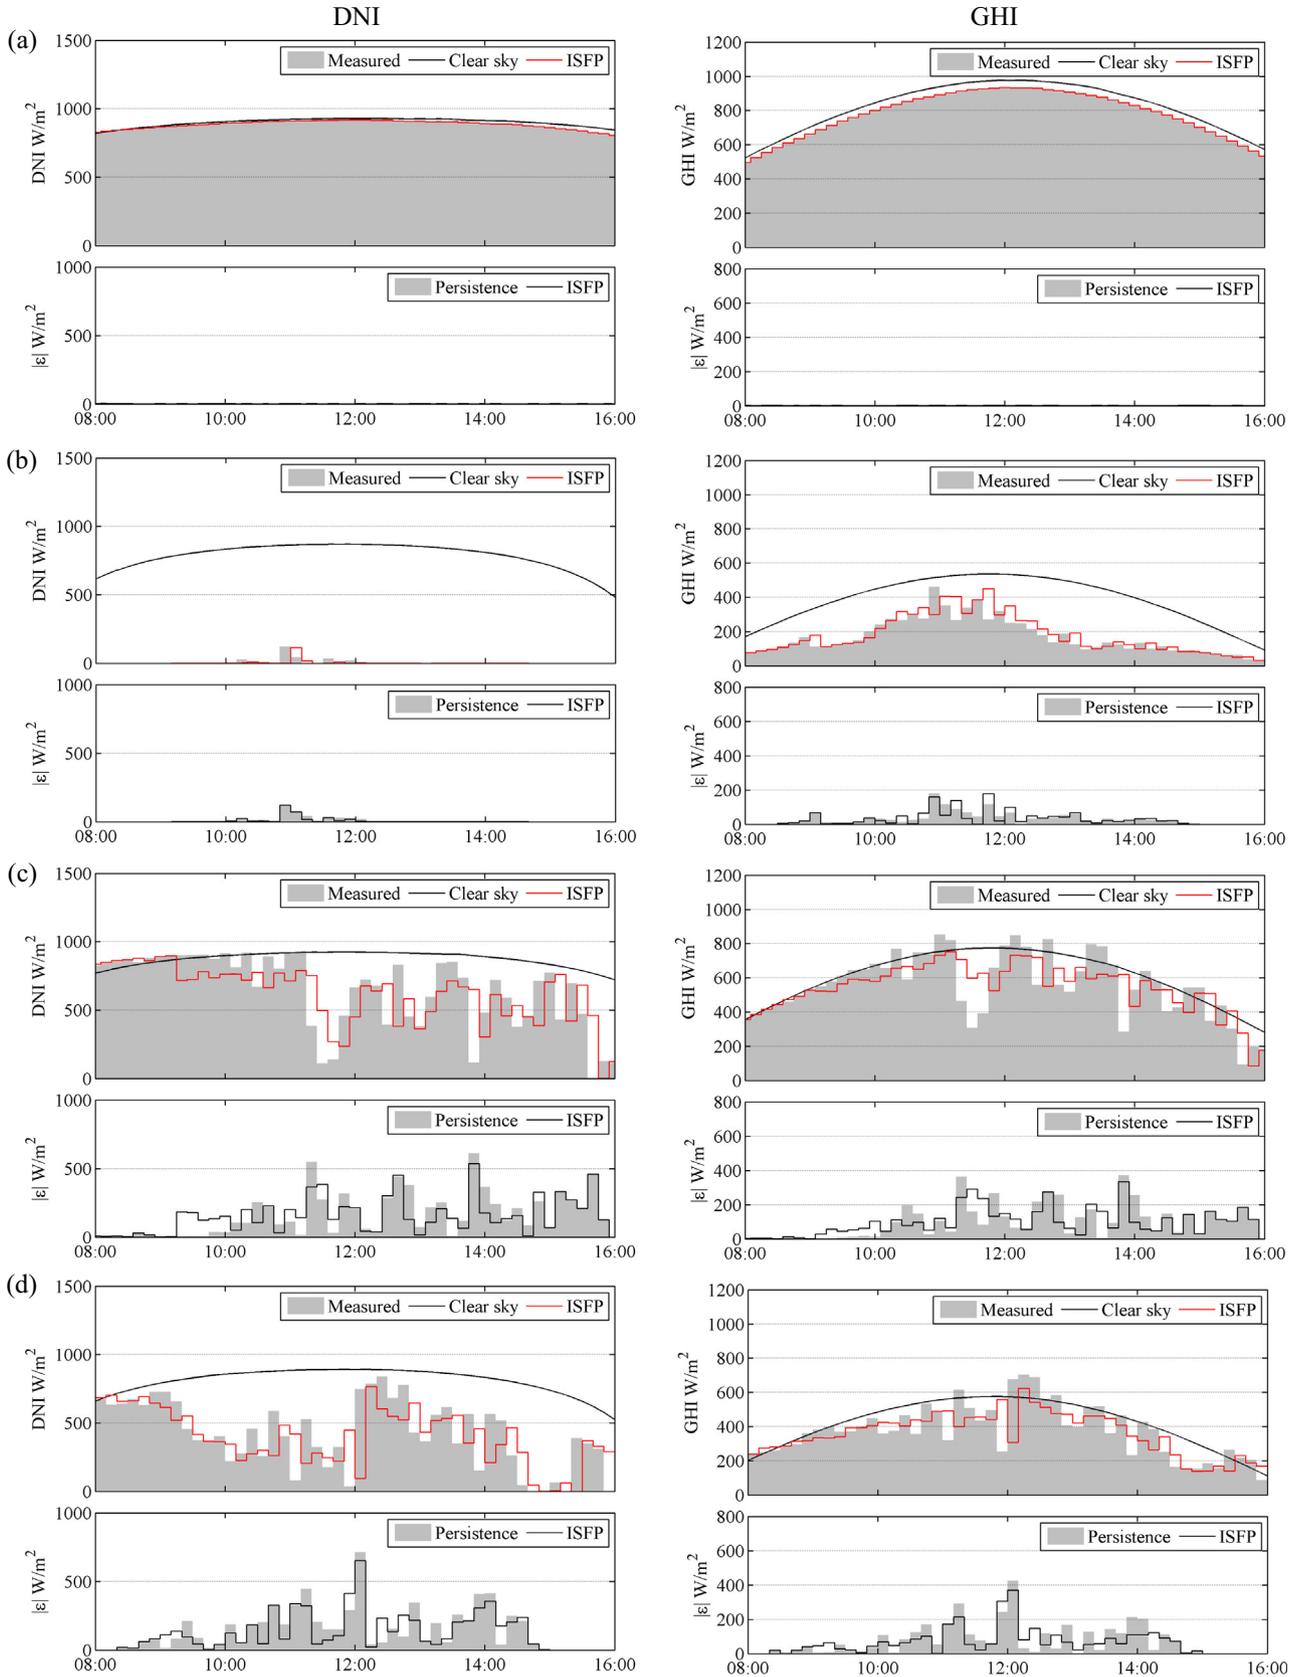| Location	Irradiance	Forecast method	MBE (W/m^2)	RMSE (W/m^2)	s (%)	k_e
Folsom	GHI	Persistence	0.7	63.6	0.0	27.6
		ISFP	3.4	57.2	10.1	23.2
	DNI	Persistence	0.9	102.1	0.0	25.8
		ISFP	1.8	90.5	11.3	18.9
San Diego	GHI	Persistence	−0.3	86.1	0.0	9.4
		ISFP	2.9	80.8	6.2	7.8
	DNI	Persistence	−1.2	151.2	0.0	10.4
		ISFP	4.7	142.1	6.0	8.6

Fig. 3. Sample time series of the 10-min forecast of 1-min averaged irradiance and absolute errors in Folsom: (a) clear period (2013-08-05), (b) overcast period (2013-11-19), (c) partly cloudy period with optically thick clouds (2013-10-06), and (d) partly cloudy period optically thin clouds (2013-11-11). The 1-min averaged data extends over the 10-min interval for ease of visualization.
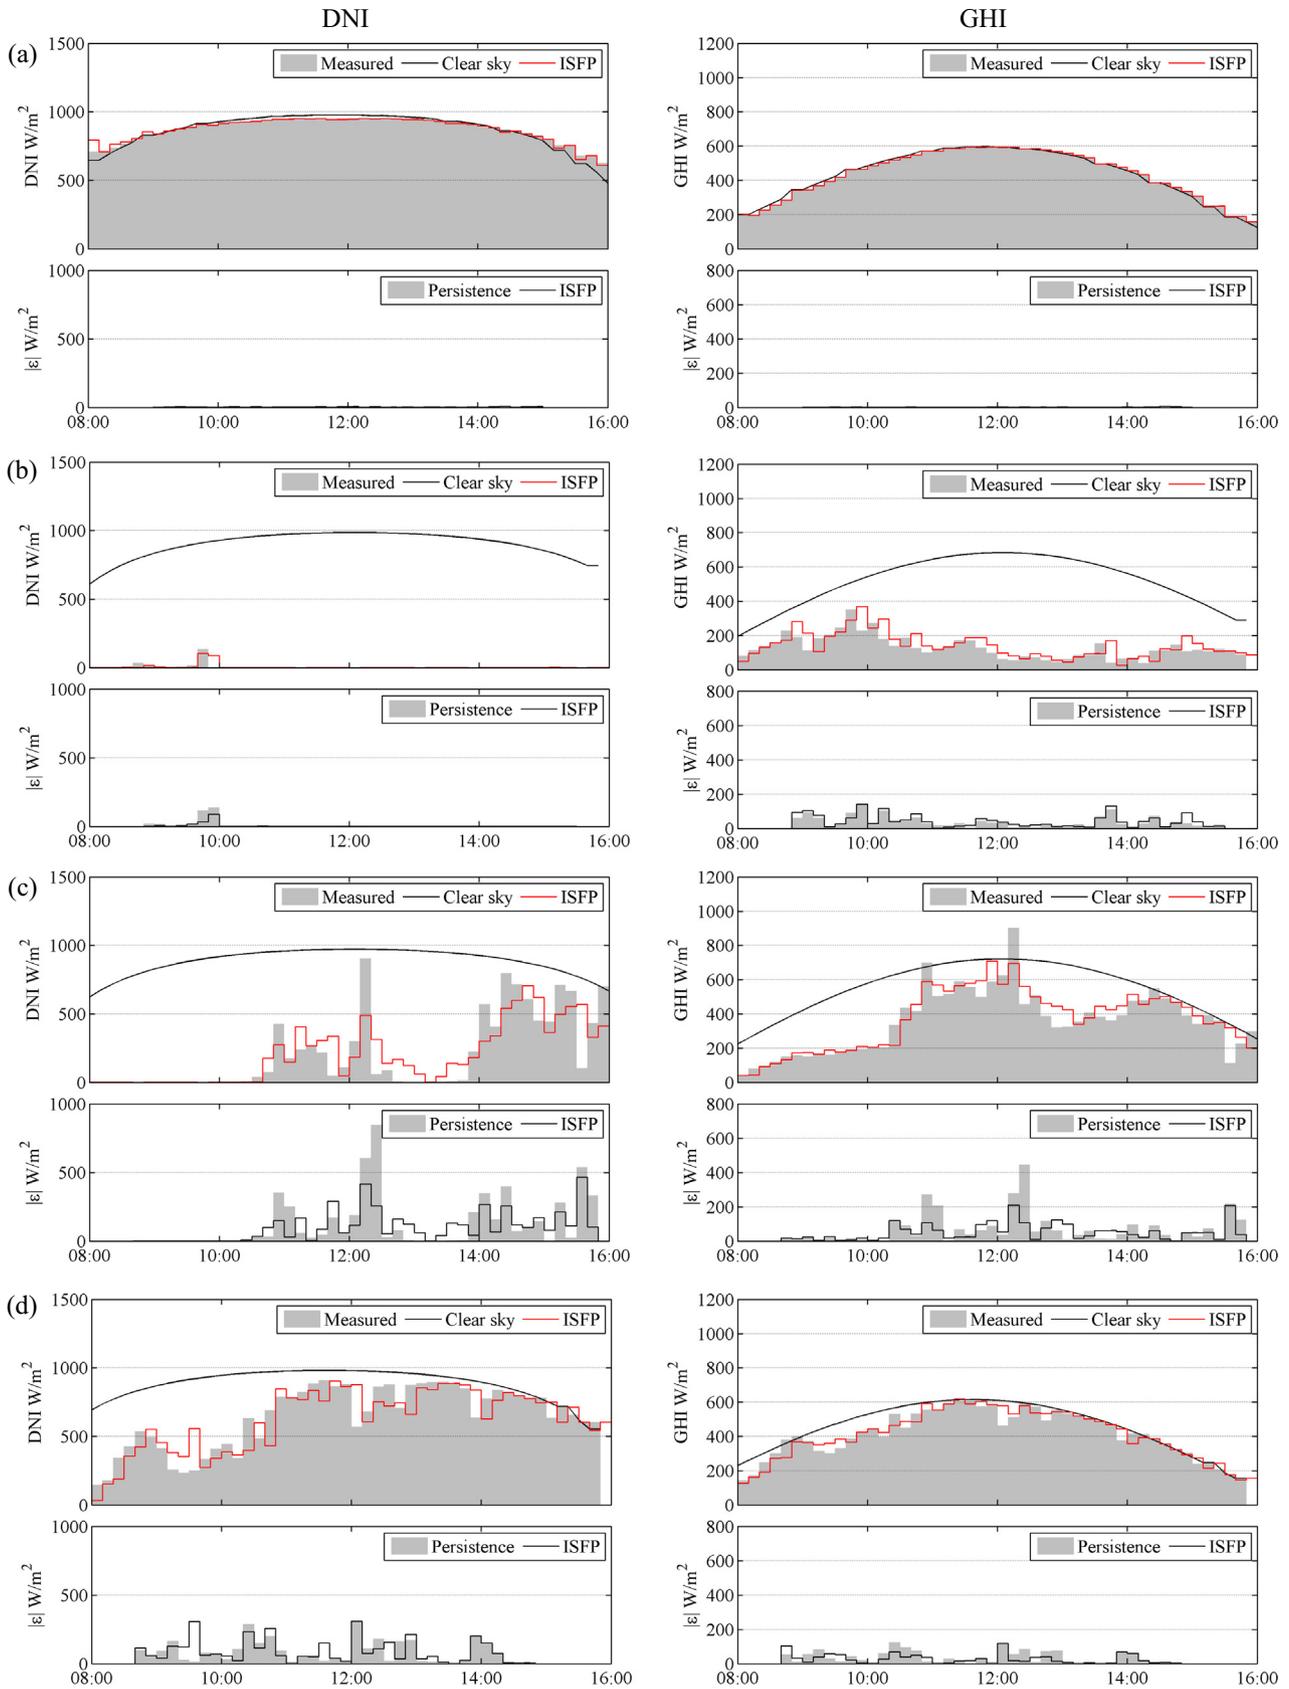
Fig. 4. Sample time series of the 10-min forecast of 1 min averaged irradiance and absolute errors in San Diego: (a) clear period (2013-12-30), (b) overcast period (2014-01-30), (c) partly cloudy period with optically thick clouds (2014-02-09), and (d) partly cloudy period optically thin clouds (2013-12-01). The 1-min averaged data extends over the 10-min interval for ease of visualization.

Table 4
Results for ramp forecasts. RDI and RMI are calculated for ramps with different ranges of non-dimensional ramp magnitudes during the analysis periods.

Location	Irradiance	Criteria	$0.1 < RM^* < 0.2$ (%)	$0.2 < RM^* < 0.3$ (%)	$0.3 < RM^* < 0.5$ (%)	$RM^* > 0.5$ (%)
Folsom	GHI	RDI	23.50	43.40	55.00	67.30
		RMI	−25.00	12.70	24.50	30.00
	DNI	RDI	28.70	43.00	45.30	72.90
		RMI	−18.00	8.20	20.00	28.10
San Diego	GHI	RDI	26.60	35.60	52.40	65.90
		RMI	−21.00	−0.10	28.10	28.50
	DNI	RDI	19.40	42.40	45.00	69.80
		RMI	−30.00	5.80	15.90	25.10

identification algorithm does not have the ability to detect the transmittance of clouds, the irradiance attenuation caused by different cloud genre are considered to be the same by ISFP. (2) Image pixels representing thin clouds usually have lower RBRs. Therefore, they are likely to be misidentified as clear pixels due to the subtraction process of CSL method, especially in the circumsolar region(e.g. Fig. 1d and h). (3) The reflection from surrounding obstacles may decrease the accuracy of irradiance measurements. (4) The errors of the clear-sky model adversely affect the quality of MLP inputs and degenerate the forecast performance, particularly for DNI forecast.

During the analysis periods, Folsom has experienced more 'clear time' than San Diego which leads to lower average irradiance variances (shown in the Table 1). Therefore, the forecasts for Folsom has smaller RMSEs, sharper peaks in error distributions, and higher kurtosis than the forecasts for San Diego. In addition, ISFP achieves significantly higher forecast skills in Folsom than in San Diego for both GHI and DNI because that the ISFP is trained with irradiance and sky image data collected in Folsom. This result shows that the ISFP performance depends on local microclimate and should be trained with irradiance and sky image data that collected locally to achieve maximum forecast skills.

4.3. Ramp prediction

We use ramp detection index (RDI), false ramp prediction (FRI) and ramp magnitude index (RMI) to assess the ISFP performance in forecasting ramps. The results are presented in Table 4 and illustrated in Figs. 5 and 6. The persistence forecast is unable to predict ramps, and its forecast errors equal to the magnitude of ramps (RDI and RMI always equal to 0). Therefore, we do not list the results of persistence forecasts in Table 4. Table 4 shows that for high-magnitude ramps ($RM^* > 0.5$), the ISFP achieves RDI >65% and RMI >25%. For moderate-magnitude ramps ($0.2 < RM^* < 0.5$), RDI ranges from 40% to 60% and RMI ranges from 0% to 30%. For low magnitude ramps $0.1 < RM^* < 0.2$), the ISFP achieves relative low RDI (ranges from 20% to 35%) and negative RMIs. Figs. 5 and 6 show that both RDI and RMI increase with the non-dimensional ramp magnitude at both locations. ISFP
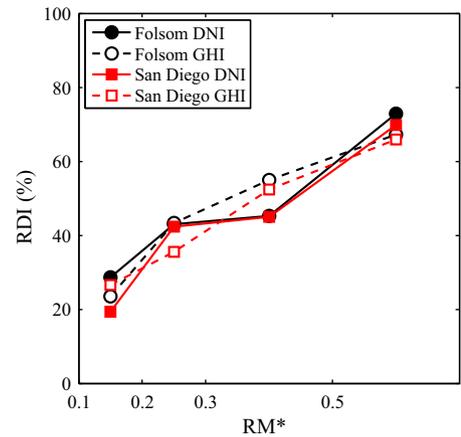

Fig. 5. RDI with respect to non-dimensional ramp magnitude for the ISFP. Round markers and square markers represent forecasts in Folsom and San Diego, respectively. RDIs are calculated for ramps within the preset non-dimensional ramp magnitude ranges (shown in Table 4).
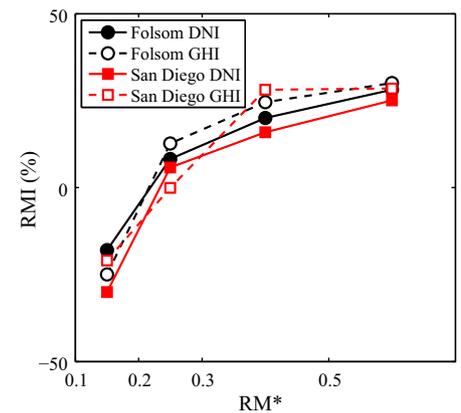

Fig. 6. RMI with respect to non-dimensional ramp magnitude for the ISFP. Round markers and square markers represent forecasts in Folsom and San Diego, respectively. RMIs are calculated for ramps within the preset non-dimensional ramp magnitude ranges (shown in Table 4).

is more accurate in forecasting ramps with high magnitude. The occurrence rate of the false DNI ramp predictions (FRI) are 2.6% and 4.2% for Folsom and San Diego, respectively. The GHI FRI are 2.3% and 4.3% for Folsom and San Diego, respectively. Most of the false ramp predictions come from partly cloudy periods (shown in Figs. 3 and 4). As discussed above, during the analysis period Fol-
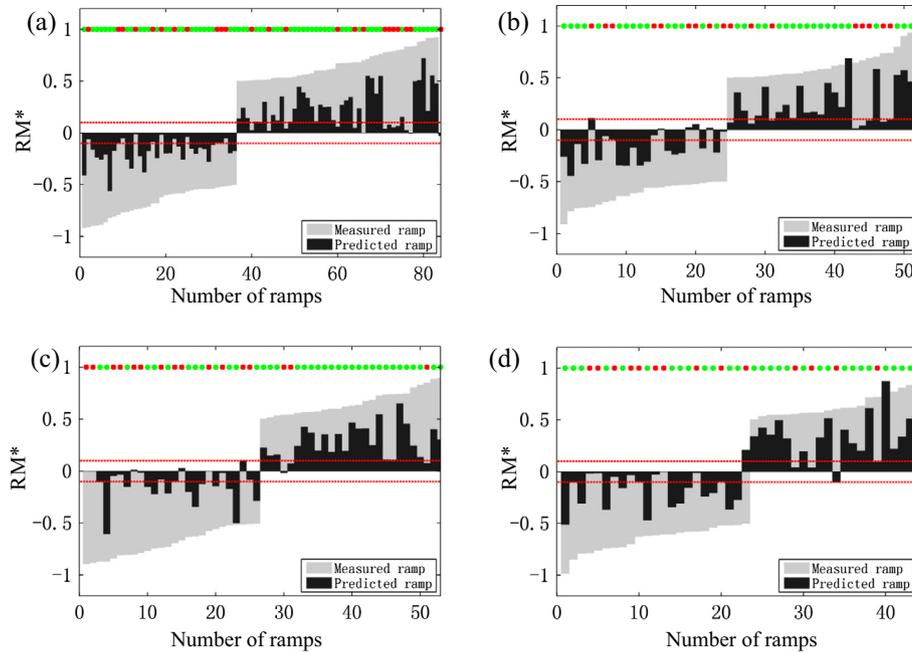
Fig. 7. Plot of the high magnitude ramps (RM* > 0.5) and corresponding ISFP predictions for (a) Folsom DNI, (b) Folsom GHI, (c) San Diego DNI, and (d) San Diego GHI. The y-axis represents the non-dimensional ramp magnitudes. The ramps are shown as bars sorted in ascended order. Negative bars indicate down ramps and positive bars indicate up ramps. The dash lines represent the ramp thresholds. The green circle and red square markers at the top of the figure represent 'hit' and 'miss', respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
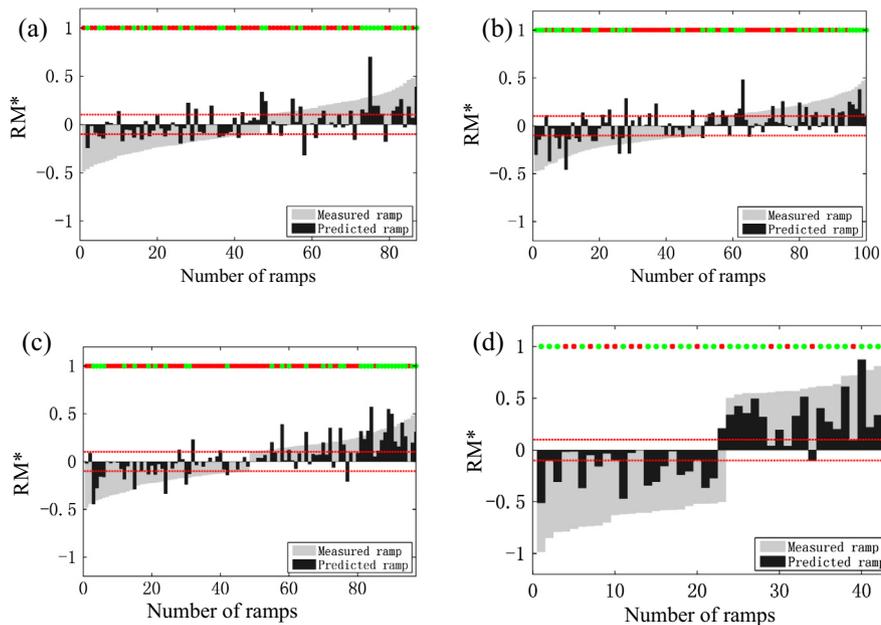

Fig. 8. Plot of the randomly selected low and moderate magnitude ramps (0.1 < RM* < 0.5) and corresponding ISFP predictions for (a) Folsom DNI, (b) Folsom GHI, (c) San Diego DNI, and (d) San Diego GHI. The y-axis represents the non-dimensional ramp magnitudes. The ramps are shown as bars sorting in ascended order. Negative bars indicate down ramps and positive bars indicate up ramps. The dash lines represent the ramp thresholds. The green circle and red square markers at the top of the figure represent 'hit' and 'miss', respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

som experiences higher fraction of clear time than San Diego. As a result, the forecast FRIs in Folsom is significantly lower than the FRIs in San Diego.

To further understand the ISFP performance in ramp forecast, we plot the high-magnitude ramps (RM* > 0.5) and corresponding ISFP predictions in Fig. 7. Fig. 7 shows that the ISFP has accuracy >65% in predicting the

occurrence of high magnitude ramps. However, the ISFP tends to underestimate the magnitude of these ramps. The underestimated ramp magnitudes are sometimes below the ramp threshold ($RM^* < 10\%$) resulting in incorrect 'no ramps' predictions. For high-magnitude ramps, most of the 'miss' cases come from these incorrect 'no ramp' predictions. Fig. 7 also shows that DNI has more high-magnitude ramps than GHI in both Folsom and San Diego during the analysis periods. This phenomenon is expected because DNI is statistically more variant than GHI.

Due to large amount of low and moderate magnitude ramps, we randomly select 20% of ramps observed in Folsom and 33% ramps observed in San Diego with RM^* ranging from 0.1 to 0.5 and plot them in Fig. 8. Three typical errors are observed in forecasting low to moderate magnitude ramps: (1) overestimation of ramp magnitude, (2) incorrect sign of ramp value, and (3) incorrect 'no ramp' predictions, especially for low magnitude ramps (Fig. 8). Both RDI and RMI decreases with the decrease of ramp magnitude. These errors are caused by two reasons: (1) the cloud classification algorithm returns binary cloud maps that unable to model the attenuation factor of clouds. Low magnitude ramps usually caused by the thin clouds with various optical path is a challenging situation for forecast. Therefore, the prediction of ramp magnitude may not accurately match the ramp events (e.g. Figs. 3d and 4d). (2) The SACI can differentiate clouds from image glare, but not very accurately in the circumsolar region (e.g. Fig. 1c, d, and h). The misidentification of cloud pixels adversely affect the accuracy of ISFP forecast, results in 'false ramp' predictions (e.g. Fig. 3d).

5. Conclusions

We proposed an MLP-enhanced, image-based, Integrated Solar Forecasting Platform (ISFP) and applied the operational forecasting platform to real-time forecasting at two locations with widely different solar microclimates. The ISFP uses the SACI method to produce binary cloud maps from sky images and then uses a grid-cloud-fraction method to extract numerical cloud indexes (CIs) from the cloud maps. Sky images are collected by fisheye dome network cameras. Thereafter, we train the ISFP with lagged measured irradiance and CIs collected from Folsom during a period of six months, and deployed the trained ISFP in real-time to perform 10 min horizon forecast at our observatories in Folsom and San Diego, California. Our operational method achieves the following forecasting skills over the reference persistence model: 10.1% for Folsom GHI, 11.3% for Folsom DNI, 6.2% for San Diego GHI, and 6.0% for San Diego DNI.

To assess the ability of the ISFP to predict ramps, we proposed three distinct metrics: (1) RDI, the percentage of ramp occurrence that is successfully predicted; (2) FRI, the percentage of false ramp predictions during no-ramp periods, and (3) RMI, the accuracy of predicting the magnitude of ramps. Both the RDI and the RMI for the ISFP are positively correlated to the non-dimensional magnitudes of ramps. For high magnitude ramps ($RM^* > 0.5$), the ISFP achieves RDIs >65% and RMIs >25%. FRIs of the ISFP are lower than 5% for both locations.

Suggestions for Future work include but not limited to: (1) Improve the ISFP algorithm and enable the ISFP to work more efficiently with higher temporal resolution for multiple horizons. (2) Further improve the accuracy of the cloud detection system. (3) Quantify the forecast uncertainty of ISFP and provide prediction intervals.

References

Bacher, P., Madsen, H., Nielsen, H.A., 2009. Online short-term solar power forecasting. Solar Energy 83, 1772–1783.

Bishop, C.M., 1994. Neural networks and their applications. Rev. Sci. Instrum. 65, 1803–1832.

Cao, J., Lin, X., 2008. Study of hourly and daily solar irradiation forecast using diagonal recurrent wavelet neural networks. Energy Convers. Manage. 49, 1396–1406.

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., Coimbra, C.F.M., 2013. Hybrid intra-hour DNI forecasts with sky image processing enhanced by stochastic learning. Solar Energy 98, 592–603.

Chu, Y., Pedro, H., Nonnenmacher, L., Inman, R.H., Liao, Z., Coimbra, C.F.M., 2014. A smart image-based cloud detection system for intra-hour solar irradiance forecasts. J. Atmos. Oceanic Technol. 31, 1995–2007.

Chu, Y., Urquhart, B., Gohari, S.M.I., Pedro, H., Kleissl, J., Coimbra, C.F.M., 2015. Short-term reforecasting of power output from a 48 MWe solar PV plant. Solar Energy 112, 68–77.

Efron, B., Gong, G., 1983. A leisurely look at the bootstrap, the jackknife, and cross-validation. Am. Statist. 37, 36–48.

Elizondo, D., Hoogenboom, G., McClendon, R.W., 1994. Development of a neural network model to predict daily solar radiation. Agric. Forest Meteorol. 71, 115–132.

Florita, A., Hodge, B., Orwig, K., 2013. Identifying wind and solar ramping events. Green Technologies Conference, 2013 IEEE. IEEE, pp. 147–152.

Geisser, S., 1993. In: Predictive Inference, vol. 55. Chapman & Hall/CRCr, New York, NY.

Ghonima, M.S., Urquhart, B., Chow, C.W., Shields, J.E., Cazorla, A., Kleissl, J., 2012. A method for cloud detection and opacity classification based on ground based sky imagery. Atmos. Measure. Techn. Discuss. 5, 4535–4569.

Hammer, A., Heinemann, D., Lorenz, E., Luckehe, B., 1999. Short-term forecasting of solar radiation: a statistical approach using satellite data. Solar Energy 67, 139–150.

Hart, E.K., Stoutenburg, E.D., Jacobson, M.Z., 2012. The potential of intermittent renewables to meet electric power demand: current methods and emerging analytical techniques. Proc. IEEE 100, 322–334.

Hassanzadeh, M., Etezadi-Amoli, M., Fadali, M.S., 2010. Practical approach for sub-hourly and hourly prediction of pv power output. In: North American Power Symposium (NAPS), pp. 1–5.

Holland, J.H., 1992. In: Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence. MIT Press.

Huang, Y., Lu, J., Liu, C., Xu, X., Wang, W., Zhou, X., 2010. Comparative study of power forecasting methods for pv stations. Power System Technology (POWERCON), 2010 International Conference. IEEE, pp. 1–6.

Ineichen, P., Perez, R., 2002. A new airmass independent formulation for the Linke turbidity coefficient. Solar Energy 73, 151–157.

Inman, R., Pedro, H., Coimbra, C.F.M., 2013. Solar forecasting methods for renewable energy integration. Progress Energy Combust. Sci. 39, 535–576.

Jain, A.K., Duin, R.P.W., Mao, J., 2000. Statistical pattern recognition: a review. IEEE Trans. Pattern Anal. Mach. Intell. 22, 4–37.

Joanes, D.N., Gill, C.A., 1998. Comparing measures of sample skewness and kurtosis. J. Royal Stat. Soc.: Series D (The Statistician) 47, 183–189.

Kalogirou, S.A., 2001. Artificial neural networks in renewable energy systems applications: a review. Renew. Sustain. Energy Rev. 5, 373–401.

Kamath, C., 2010. Understanding wind ramp events through analysis of historical data. Transmission and Distribution Conference and Exposition, 2010 IEEE PES. IEEE, pp. 1–6.

Kenny, J.F., 1961. Mathematics of Statistics. Nostrand.

Kohavi, R., 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. In: International Joint Conference on Artificial Intelligence, pp. 1137–1145.

Lave, M., Kleissl, J., Stein, J.S., 2012. A wavelet-based variability model (wvm) for solar pv power plants. IEEE Trans. Sustain. Energy 4, 501–509.

Lendasse, A., Wertz, V., Verleysen, M., 2003. Model selection with cross-validations and bootstraps-application to time series prediction with RBFN models. Artif. Neural Networks Neural Informat. Process. ICANN/ICONIP 2714, 573–580.

Li, C.H., Lee, C.K., 1993. Minimum cross entropy thresholding. Pattern Recogn. 26, 617–625.

Li, C.H., Tam, P.K.S., 1998. An iterative algorithm for minimum cross entropy thresholding. Pattern Recogn. Lett. 19, 771–776.

Li, Q., Lu, W., Yang, J., 2011. A hybrid thresholding algorithm for cloud detection on ground-based color images. J. Atmos. Oceanic Technol. 28, 1286–1296.

Li, Y., Luan, R., Niu, J., 2008. Forecast of power generation for grid-connected photovoltaic system based on grey model and markov chain. In: ICIEA 2008 3rd IEEE Conference, pp. 1729–1733.

Long, C.N., Ackerman, T.P., 2000. Identification of clear skies from broadband pyranometer measurements and calculation of downwelling shortwave cloud effects. J. Geophys. Res. 105, 15609–15626.

Marquez, R., Coimbra, C.F.M., 2011. Forecasting of global and direct solar irradiance using stochastic learning methods, ground experiments and the nws database. Solar Energy 85, 746–756.

Marquez, R., Coimbra, C.F.M., 2013. Intra-hour DNI forecasting methodology based on cloud tracking image analysis. Solar Energy 91, 327–336.

Marquez, R., Gueorguiev, V.G., Coimbra, C.F.M., 2013a. Forecasting of global horizontal irradiance using sky cover indices. ASME J. Solar Energy Eng. 135, 0110171–0110175.

Marquez, R., Pedro, H.T.C., Coimbra, C.F.M., 2013b. Hybrid solar forecasting method uses satellite imaging and ground telemetry as inputs to anns. Solar Energy 92, 176–188.

Martín, L., Zarzalejo, L.F., Polo, J., Navarro, A., Marchante, R., Cony, M., 2010. Prediction of global solar irradiance based on time series analysis: application to solar thermal power plants energy production planning. Solar Energy 84, 1772–1781.

McLachlan, G.J., Do, K.A., Ambroise, C., 2004. In: Analyzing Micro-array Gene Expression Data, vol. 422. Wiley-Interscience, Hoboken, NJ.

Mellit, A., Eleuch, H., Benghanem, M., Elaoun, C., Pavan, A.M., 2010. An adaptive model for predicting of global, direct and diffuse hourly solar irradiance. Energy Convers. Manage. 51, 771–782.

Mellit, A., Kalogirou, S.A., 2008. Artificial intelligence techniques for photovoltaic applications: a review. Progress Energy Combust. Sci. 34, 574–632.

Mellit, A., Pavan, A.M., 2010. A 24-h forecast of solar irradiance using artificial neural network: application for performance prediction of a grid-connected pv plant at trieste, italy. Solar Energy 84, 807–821.

Mori, N., Chang, K.A., 2003. Introduction to mpiv. User manual and program, <http://sauron.civil.eng.osaka-cu.ac.jp/~mori/softwares/mpiv>.

Nonnenmacher, L., Coimbra, C.F.M., 2014. Streamline-based method for intra-day solar forecasting through remote sensing. Solar Energy 108, 447–459.

Nonnenmacher, L., Kaur, A., Coimbra, C.F.M., 2014. Verification of the suny direct normal irradiance model with ground measurements. Solar Energy 99, 246–258.

Otsu, N., 1979. A threshold selection method from gray level histograms. IEEE Trans. Syst. Man. Cybern. 9, 62–66.

Pai, P.F., Hong, W.C., 2005. Forecasting regional electricity load based on recurrent support vector machines with genetic algorithms. Electric Power Syst. Res. 74, 417–425.

Pedro, H.T.C., Coimbra, C.F.M., 2012. Assessment of forecasting techniques for solar power production with no exogenous inputs. Solar Energy 86, 2017–2028.

Perez, R., Kivalov, S., Schlemmer, J., Hemker, K., Renne, D., Hoff, T.E., 2010. Validation of short and medium term operational solar radiation forecasts in the US. Solar Energy 84, 2161–2172.

Quesada-Ruiz, S., Chu, Y., Tovar-Pescador, J., Pedro, H.T.C., Coimbra, C.F.M., 2014. Cloud-tracking methodology for intra-hour DNI forecasting. Solar Energy 102, 267–275.

Reno, M.J., Hansen, C.W., Stein, J.S., 2012. Global Horizontal Irradiance Clear Sky Models: Implementation and Analysis. Sandia National Laboratories.

Singh, G.K., 2013. Solar power generation by pv (photovoltaic) technology: a review. Energy 53, 1–13.

Urquhart, B., Ghonima, M., Nguyen, D., Kurtz, B., Kleissl, J., 2013. Solar Energy Forecasting and Resource Assessment. Elsevier.

Younes, S., Muneer, T., 2007. Clear-sky classification procedures and models using a world-wide data-base. Appl. Energy 84, 623–645.

Zhang, J., Hodge, B., Florita, A., Lu, S., Hamann, H.F., Banunarayanan, V., 2013. Metrics for evaluating the accuracy of solar power forecasting. Technical Report. National Renewable Energy Laboratory, NREL/CP-5500-60142.

Zheng, H., Kusiak, A., 2009. Prediction of wind farm power ramp rates: a data-mining approach. J. Solar Energy Eng. 131, 0310111–0310118.

Zhang, J., Florita, A., Hodge, B., Lu, S., Hamann, H.F., Banunarayanan, A., Brockway, A.M., 2015. A suite of metrics for assessing the performance of solar power forecasting. Solar Energy 111, 157–175.