Irradiance Forecasts based on an Irradiance Monitoring Network, Cloud Motion, and Spatial Averaging

Antonio T. Lorenzo^{a,*}, William F. Holmgren^b, Alexander D. Cronin^{c,a}

^aUniversity of Arizona, College of Optical Sciences, 1630 E. University Blvd., Tucson, AZ 85721, United States

^bUniversity of Arizona, Department of Atmospheric Sciences, 1118 E. 4th Street, Tucson, AZ 85721, United States

^cUniversity of Arizona, Department of Physics, 1118 E. 4th Street, Tucson, AZ 85721, United States

Abstract

We describe and evaluate forecasts of solar irradiance using real-time measurements from a network of irradiance sensors. A forecast method using cloud motion vectors obtained from a numerical weather model shows significant skill over a standard persistence model for forecast horizons from 1 min to over 2 h, although the skill metric may be misleading. To explain this finding, we define and compare several different persistence methods, including persistence methods informed by an instantaneous spatial average of irradiance sensor output and persistence forecasts informed by a time-average of recent irradiance measurements. We show that spatial- or temporal-averaging reduces the forecast RMS errors primarily because these forecasts are smoother (have smaller variance). We use a Taylor diagram, which shows correlation, RMSE, and variance, to more accurately compare several different types of forecasts. Using this diagram, we show that forecasts using the network of sensors have meaningful skill up to 30 min time horizons after which the skill is primarily due to smoothing.

Keywords: Solar forecasting, solar irradiance, sensor network

1. Introduction

The intermittency of solar power causes a cost to utilities and, ultimately, rate payers (Joskow, 2011). Solar power forecasts (Kleissl, 2013; Inman et al., 2013) may reduce these costs by enabling utilities to manage the variability of solar power in a number of ways. For example, forecasts can be used in conjunction with battery storage systems to control

ramp-rates or provide frequency support (Hill et al., 2012; Cormode, 2015). Additionally, forecasts will provide utility grid operators with a prediction of the expected photovoltaic (PV) output so they can more efficiently schedule backup generators.

A number of different techniques are used to forecast global horizontal irradiance (GHI). For forecast horizons in the intra-minute to a few minute range, techniques with input data from several ground sensors are often used (Achleitner et al., 2014; Elsinga and van Sark, 2014; Yang et al., 2015; Lipperheide et al., 2015).

*Corresponding author

Email address: atlorenzo@email.arizona.edu

(Antonio T. Lorenzo)

For longer (intra-hour) forecast horizons, methods based on irradiance sensor networks (Lonij et al., 2013), machine learning techniques (Chu et al., 2015b), and sky imagers (Yang et al., 2014; Chu et al., 2015a) are being actively studied. Satellite image based forecasts are useful for 1-h to many hours in advance (Perez et al., 2010; Bilonis et al., 2014). For time horizons from several hours to up to a week in advance, numerical weather models often give the best predictions (Mathiesen and Kleissl, 2011; Diagne et al., 2014; Perez et al., 2013). Combinations of techniques are also being studied to extend the useful time horizon of a forecast (Marquez et al., 2013; Lauret et al., 2014).

Networks of irradiance sensors overcome some challenges typically associated with sky imagers or satellite images. For example, data from networks of irradiance sensors do not have the issue of converting pixel brightness to irradiance as sky imagers and satellite image methods have. Sky imagers and satellite images have the additional challenge of estimating cloud height to correctly project irradiance at cloud height to a location on the ground.

In this paper, we describe GHI forecasts that utilize a network of sensors placed throughout Tucson, AZ for April, May, and June 2014. The ideas behind this work are similar to those of Lonij et al. (2013), however the data sources and implementation are different. The rooftop PV network in Lonij et al. (2013) was limited to historical reports of 15 min average power, whereas the irradiance sensors used in the present research report 1 s resolution data with 1 min latency. This allows us to make higher resolution and, as we will see, more accurate forecasts.

We will show that our sensor network based forecasting method has significant skill when compared to a clear-sky index persistence forecast from 1 min to beyond 2 h time horizons. While the limited area and density of the network likely limits the skill and forecast horizon of our network-based forecasting method, the geographic diversity of measurements provide several advantages including improved persistence estimations. We will also explore why the forecasts exhibit such significant skill and explain this result is due to smoothing after 30 min forecast horizons.

First, we describe our network of irradiance sensors. Then, we describe how we use the network to make forecasts. A discussion of different types of persistence forecasts follows. Finally, we present and discuss our results and offer a concluding summary.

2. Irradiance Sensor Network Forecasts

Our forecasting method relies on a network of sensors that sample the global horizontal irradiance at a number of locations. Our network consists of 12 irradiance sensors we developed, plus three rooftop PV power systems and one calibrated, commercial sensor. The calibrated sensor is part of a National Renewable Energy Laboratory (NREL) Solar Resource and Meteorological Assessment Project (SOLRMAP) site at the Univ. of Arizona (Wilcox and Andreas, 2010). Converting the data to clear-sky indices using an expected clear-sky profile for each sensor allows us to combine sensors that measure different quantities to make forecasts. These sensors are distributed throughout Tucson as shown

in Fig. 1. The irradiance sensors we developed collect 1 s data and transmit it to a database every minute via cellular data networks (Lorenzo et al., 2014). Some use commercial pyranometers while others use photodiodes. Since we use clear-sky indices with data driven clear-sky profiles, the absolute error of the sensor is not a concern. However, the sensor used to evaluate the forecasted irradiance is a commercial sensor (Apogee SP-212) and agrees with the calibrated sensor to within 2% on average on clear days. The data was plotted for each day and for each sensor and verified by eye to provide some measure of quality control. See Lorenzo et al. (2015) for access to the dataset that was used in this study.

The first step in making our forecasts is to convert irradiance and PV power data to clear-sky index data. The clear-sky index for a sensor n at time t is defined as

$$k_n(t) = \frac{y_n(t)}{y_n^{clr}(t)} \quad (1)$$

where $y_n(t)$ is the measured data and $y_n^{clr}(t)$ is the clear-sky expectation. Clear-sky expectations for each sensor are generated by fitting the measured data on a clear day in the recent past. An advantage of using this data-driven method of generating clear-sky expectations rather than a clear-sky model, such as the REST2 model (Gueymard, 2008) or Ineichen model (Ineichen and Perez, 2002), is that the data-driven method inherently accounts for sensor orientation, permanent obstacles, and sensor calibration errors. Furthermore, because our forecasting method relies on forecasting clear-sky index and then converting back to irradiance as a final step, the changes in atmospheric conditions

Figure 1: Map of irradiance sensors used for this study in Tucson, AZ. The red star indicates the position of the sensor that was used to evaluate forecasts in Section 5. The sensor was chosen because of its proximity to 25 MW of installed PV power in and around the University of Arizona Science and Technology Park Solar Zone. The forecast area extends from 31.83° N to 32.28° N and 110.7° W to 111.15° W.

between clear-sky days are not a major source of error in our final forecasts.

Next, we use the clear-sky indices to interpolate the scattered data onto a 1400 km² clear-sky index map for the Tucson region. To generate a clear-sky index map, we first create a 0.001° grid (with grid points approximately every 100 m) and add the sparse data from our network to the grid. We then set points along each boundary of the grid to the average of the clear-sky index values obtained from all the sensors. As we will see in Sections 4.4 and 5, this boundary condition helps to maintain forecast skill at longer time horizons.

Next, we fill all points in the grid with interpo-

lated values as shown in Fig. 2. We chose to use multiquadric interpolation because it performs well at interpolating scattered geospatial data (Franke, 1982; Nuss and Titley, 1994), and it was more robust with our sparse data. We did not use a kriging method, even though kriging is often used to interpolate geospatial data, because we lack sufficient data to adequately estimate the variogram (Webster and Oliver, 1993; Sirayanone, 1988). We also explored inverse distance weighted interpolation, but found the output to be similar to multiquadric interpolation with insignificant differences in forecast errors (typically $< 3 \text{ W/m}^2$ difference in the root mean square error for all forecast horizons).

Then, we translate this interpolated clear-sky map a distance determined by the cloud motion vectors (which may vary in time). The translation in the x direction, with the y translation being analogous, is given by

$$\Delta x(t_i, t) = \int_{t_i}^t v_x(t') dt', \quad (2)$$

where t_i is the time at which the forecast is being made, $t - t_i$ is the forecast horizon, and $v_x(t)$ is the x component of the time-varying cloud motion vector. Any grid points that are missing data after the translation are filled with the average clear-sky index for all the sensors. Figure 2 shows an example of an interpolated clear-sky index map and a map that has been shifted along the estimated cloud motion vector. Finally, we sample from this translated map at the desired forecast locations to obtain a forecasted clear-sky index which can be multiplied by the clear-sky expectation for that location to obtain an irradiance forecast. As we will discuss, for sufficiently long forecast horizons this procedure

makes our network based forecasts indistinguishable from spatially-average persistence forecasts.

Forecasts out to 2 h in advance with 1 min time resolutions were made every 1 min for this analysis. As an example, one hour's worth of 5 min ahead forecasts along with measurements are shown in Fig. 3. This time-series is a composite showing snapshots (individual points) from 120 different forecasts that were each made 5 min in advance on a rolling basis. Concatenating points from different forecasts this way for several months lets us evaluate errors for forecasts with a 5 min horizon. Furthermore, updating forecasts every 1 min is valuable since 1 s data is constantly streaming into our database and each forecast has some new information that will likely improve the prediction for a specific time in the future. Later, we discuss errors as function of forecast horizon.

For estimating cloud motion velocity vector components (v_x, v_y), several techniques have been discussed including sensor correlations (Fung et al., 2014; Bosch et al., 2013), predictions from NWP (Lave and Kleissl, 2013; Lonij et al., 2013), analysis of aircraft communications addressing and reporting system (ACARS) or rawinsonde data, scaling of measured ground velocity, analysis of sky camera images (Urquhart et al., 2013), and analysis of satellite images (Hammer et al., 1999). For our analysis, we used modeled soundings (atmospheric temperature and dewpoint as a function of altitude/pressure) from the Weather Research and Forecasting (WRF) model run by the Univ. of Arizona, Dept. of Atmospheric Sciences (Leuthold, 2015). First, we compute a profile of relative humidity as a function of altitude averaged over the

Figure 2: An example interpolated map of clear-sky index on 5/19/14 near noon is shown in a). Using the estimated cloud motion vectors this map is shifted according to desired forecast horizon as shown in b). Then, samples from this shifted map are taken to as the forecasted clear-sky index for a particular location. The white space at bottom and left of b) is filled in with the average clear-sky index of all sensors at the time the forecast is generated. The red star indicates the sensor that was used to evaluate forecasts.

Figure 3: An example of a 5 min ahead network forecast compared to measured data. Forecasts were generated every 1 min and the forecast for 5 min in the future is shown. The forecast and measurements at 12:00 show excellent agreement. For reference, the MAE for this entire period is 105 W/m² and the RMSE is 140 W/m², and for 11:45 to 12:00 the MAE is 68 W/m² and the RMSE is 82 W/m².

Tucson area from the WRF model. To estimate the cloud motion vectors, we find the altitude at which relative humidity is greatest (dashed line in Fig. 4), similar to Lave and Kleissl (2013). We then find all nearby heights that have a relative humidity that is within 90% of the maximum (shaded area in Fig. 4). The wind speed and direction is then averaged for these altitudes and over the entire Tucson area to provide an estimated cloud motion vector. A new cloud motion vector is estimated in this way from each hourly output of the WRF model and then interpolated to 1 min time resolution. This simple estimation method has a number of limitations including only recognizing a single cloud layer and possibly selecting the wrong layer of the atmosphere i.e. one in which there are no clouds. This cloud motion estimation method along with

Figure 4: Example vertical profiles of relative humidity and wind speed made by a numerical weather model on 5/19/14 at noon. To find the altitude at which clouds are most likely to form, we find the height with the greatest relative humidity (red dashed line). The winds at this height and heights within 90% of the maximum relative humidity (red shaded area) are averaged to produce an estimate of the cloud motion vector.

the modest size and density of our network likely limits the overall accuracy of the network based forecasts presented here. Still, this network based method produces forecasts with lower errors than several standard persistence methods, as we discuss next.

3. Error Metrics

We assessed the accuracy of forecasts using standard error metrics that are defined in Zhang et al. (2015). Each error metric is computed for forecast horizons, FH , ranging from 1 min to 30 min ($FH = 0, 1, \dots, 30$) by comparing forecasts, $y^{FH}(t_i)$,

to subsequent instantaneous measurements, $y(t_i)$, of a single irradiance sensor. Errors were only computed when the solar zenith angle was less than 75 degrees. Unless otherwise noted, only the 46 cloudy days in the study period were used to calculate error metrics and each metric is computed over this entire cloudy data set. Data and forecasts for a sensor (star in Fig. 1) in the middle of the network and near many large PV installations were used. Comparisons are always made with an instantaneous measurement, not averaged data, even when the forecast uses averaging.

In addition to root-mean squared error (RMSE) and mean absolute error (MAE), we also compute the centered root-mean squared error (CRMSE) for irradiance

$$\text{CRMSE}(FH) = \left(\frac{1}{N} \sum_{i=1}^N [(y^{FH}(t_i) - \bar{y}^{FH}) - (y(t_i) - \bar{y})]^2 \right)^{1/2}, \quad (3)$$

where an overbar indicates the sample mean of the quantity (Taylor, 2001). The CRMSE removes forecast bias and will become important later.

We also compute errors for forecasted clear-sky indices. This is valuable because, as opposed to irradiance, clear-sky index errors are not weighted based on the position of the sun in the sky.

We also define *relative* metrics in terms of clear-sky indices in order to present errors in percentages. The relative RMSE is

$$\text{rRMSE}(FH) = \bar{k}^{-1} \left(\frac{1}{N} \sum_{i=1}^N (k^{FH}(t_i) - k(t_i))^2 \right)^{1/2}. \quad (4)$$

Relative MAE is similarly defined as

$$\text{rMAE}(FH) = \bar{k}^{-1} \frac{1}{N} \sum_{i=1}^N |k^{FH}(t_i) - k(t_i)|. \quad (5)$$

Following the method of Marquez and Coimbra (2012), we can approximate forecast skill s as

$$s(FH) \approx 1 - \frac{\text{RMSE}(FH)}{\text{RMSE}_p(FH)} \quad (6)$$

where RMSE_p is the RMSE for a clear-sky persistence forecast, described in Section 4.2. To estimate the average skill over many days, the ratio $\frac{\text{RMSE}}{\text{RMSE}_p}$ is estimated by the slope of the regression fit of daily RMSE vs RMSE_p . The average skill is then $\langle s \rangle = 1 - \text{slope}$. Examples of these plots and regressions are presented in Fig. 11.

4. Persistence Forecasts

Persistence forecasts are the simplest type of forecast to implement and are often the most accurate at very short time horizons, making them a standard to compare with other methods. In this section we describe and compare the persistence forecasts we use for irradiance forecasting.

Before describing the various types of persistence, we first define the terminology we will use. The measured quantity of sensor n (e.g. irradiance) at time t will be denoted by $y_n(t)$. The forecast of sensor n at some time $t + FH$ in the future will be denoted by $y_n^*(t + FH)$. As mentioned in Sec. 3, we call FH the forecast horizon. The clear-sky expectation for a particular sensor will be denoted y_n^{clr} and the value of the clear-sky expectation at time t is $y_n^{clr}(t)$.

4.1. Measurement Persistence

We call one of the simplest persistence methods “measurement persistence.” A measurement persistence forecast simply assumes that the irradiance at a future time will be the same as it is at the current time. Measurement persistence is defined by

$$y_n^*(t + FH) = y_n(t). \quad (7)$$

This type of persistence is useful for short time horizons, but it does not account for the diurnal cycle of irradiance due to changing solar position and this leads to large errors at longer time horizons as shown in Fig 5.

4.2. Clear-Sky Index Persistence

In this method, the clear-sky index is calculated at the current time and persisted into the future. A forecast of irradiance is obtained by multiplying this clear-sky index by the value of the clear-sky expectation at the forecast time. The equation for clear-sky index persistence is

$$y_n^*(t + FH) = \frac{y_n^{clr}(t + FH)}{y_n^{clr}(t)} \times y_n(t). \quad (8)$$

This method performs better than measurement persistence because it takes into account the diurnal cycle of irradiance, but it does require that a clear-sky expectation for the sensor, $y_n^{clr}(t)$, be known or modeled appropriately.

4.3. Time-Averaged Persistence

At time horizons greater than a few minutes, it can be beneficial to first average the measured clear-sky index over some time period defined by N time steps, each with period Δt , ending at some past

time t_0 . This average clear-sky index is then multiplied by the clear-sky expectation of the target sensor to compute a forecast. Time-averaged persistence is thus computed as

$$y_n^*(t + FH) = y_n^{clr}(t + FH) \times \frac{1}{N} \sum_{i=0}^{N-1} \frac{y_n(t - t_0 - i\Delta t)}{y_n^{clr}(t - t_0 - i\Delta t)}. \quad (9)$$

Often, a rolling averaged is used so $t_0 = 0$, Δt is the time resolution of the measured data, and N is chosen so $(N - 1)\Delta t$ gives the desired averaging time. The total averaging time does not limit the frequency with which forecasts can be made. For example, a 5 min rolling average persistence can be recomputed every 1 min and still provide a useful forecast since new data is incorporated every time a forecast is made. An example of time-averaged persistence error with different averaging times using a rolling average is shown in Fig. 6.

4.4. Spatially-Averaged Persistence

If multiple measurements of irradiance are available in an area, one can make a persistence forecast based on the average clear-sky index of all the sensors. We refer to this method as spatially-averaged persistence. To make these forecasts, the measurements of multiple sensors are first converted to clear-sky indices using clear-sky expectations for each sensor. Then, these clear-sky indices are averaged together. This average clear-sky index is then multiplied by the clear-sky expectation of the target sensor to produce a forecast for that sensor. Using N sensors, the spatially-averaged persistence for sensor n is

$$y_n^*(t + FH) = y_n^{clr}(t + FH) \times \frac{1}{N} \sum_{m=1}^N \frac{y_m(t)}{y_m^{clr}(t)}. \quad (10)$$

This method does not perform as well as clear-sky index persistence or measurement persistence at time horizons under a few minutes, as shown in Fig. 5, but it is more accurate (according to the RMSE metric) than other persistence methods discussed here at longer (2 - 30 min) forecast horizons.

One could also imagine replacing the simple mean in Eq. (10) with a weighted mean by, for example, using the lasso (Yang et al., 2015) or some other shrinkage and selection method. Time and spatial averaging can also be combined as discussed in Sec. 5.1

5. Results

We now present the results of the network and persistence forecasts using metrics defined in Zhang et al. (2015) and Sec. 3 for the study period of April, May, and June 2014. First, we evaluate persistence forecasts. Then, we study network forecast errors in depth. Finally, we compare network forecasts to other irradiance forecasting methods.

5.1. Persistence Forecast Results

Root-mean squared errors from the four types of persistence forecasts described above are plotted in Fig. 5. We see that for the 46 cloudy days we studied in Tucson, AZ., the two types of input averaging, spatial and temporal, both improve forecasts compared to clear-sky index persistence after time horizons of a few minutes. The crossover time depends on the weather. As expected, clear-sky index persistence performs better than measurement persistence because it accounts for the diurnal cycle.

Though Fig. 5 shows spatially-averaged persistence outperforming time-averaged persistence, the

Figure 5: Comparison of different types of persistence forecasts. RMSE, plotted as a function of forecast horizons, was computed for each type of forecast using data from the 46 cloudy days as described in Sec. 3. Spatially-averaged persistence has the lowest RMSE for all but the very shortest forecast horizons.

averaging time and number of sensors averaged can change these curves significantly. Figures 6 and 7 show various averaging times and number of sensors in the average, respectively. We see that longer averaging times reduce errors at time horizons greater than 5 min but are worse at shorter time horizons. The common auto-regressive moving average (ARMA) model similarly weights previous values and/or errors to produce a forecast. We also see that adding more sensors to a spatially-averaged persistence reduces errors except at time horizons shorter than a few minutes.

One explanation for our finding that spatially-averaged persistence performs better than time-averaged persistence is related to the number of dimensions in each average. Using kinematics ($x = vt$) we can map the time series $y_i(t)$ onto a one-dimensional transect in space downwind from the sensor. In comparison, the spatial average uses data from locations that are distributed in two dimen-

Figure 6: Comparison of time-averaged persistence forecasts with different averaging times. The averages shown are made via a rolling average ($t_0 = 1$) with $\Delta t = 1$ s and N adjusted for each curve to give the appropriate total averaging time as described in Sec. 4.3. Longer time averages reduce errors at longer time horizons.

sions including some locations that are upwind of the location of interest. By averaging over two dimensions, not one, spatial average persistence effectively uses more independent samples of the cloud field. This theory assumes that all sensors are subject to the same cloud field, which is reasonable for the size of our network.

When we average the input data over both space and time, as shown as the green line in Fig. 8, we find the RMSE is lower at longer time horizons.

5.2. Network Forecast Results

We now compare our network forecasts to a clear sky ($k_n^*(t) = 1$) forecast, measurement persistence, clear-sky index persistence, and spatially averaged persistence (using the same 16 sensors which were used to make the network forecast). Figure 9 shows the MAE for these methods for only cloudy days while Fig. 10 shows the RMSE. Plots of CRMSE show similar trends. Note that network forecasts

Figure 7: Comparison of spatially-average persistence forecasts with a varying number of sensors averaged. Adding more sensors to the spatial average improves the forecast RMSE.

have nonzero error at zero forecast horizons because of the smoothing applied when making the interpolated clear-sky index map and due to limiting the maximum forecasted clear-sky index to 1.25. We see that network forecasts have lower MAE than other methods for time horizons from 1 min to 30 min. We only graph up to 30 min forecast horizons because the 30 min to 2 h errors are similar and uninteresting. Figure 10 shows that the network forecasts have lower RMS errors than the other methods at forecast horizons less than about 4 min and then have slightly higher RMSE values than spatially-averaged persistence. This difference between RMSE and MAE suggests that network forecasts have fewer small errors but more large errors than spatially-averaged persistence forecasts. For completeness, we also present error metrics for all 91 days in the study period in Appendix A. Clear days show similar trends but with smaller errors which lowers the 91 day average RMSE by 40-50% depending on the time horizon.

Figure 8: Comparison of a persistence forecast made by first averaging over space and then averaging over time (green line) to other persistence methods. Averaging in time and space marginally improves forecasts at longer time horizons.

Figure 9: MAE of many types of forecasts averaged over 46 cloudy days. Clear sky refers to a forecast where one assumes the sky is always clear ($k_n^*(t) = 1$). Network forecasts have the lowest MAE at all time horizons shown.

We also compute forecast skill as defined by Marquez and Coimbra (2012). Figure 11 illustrates the regressions used to calculate the average skill of our forecasts. At low clear-sky index persistence RMSE values (e.g. clear days), we see that the skill is negative (network RMSE > clear-sky index persistence RMSE). For days with larger clear-sky index persistence RMSE values, we see that our network forecasts have positive skill. The average skill found from regressions, typically 20%, is plotted in Fig. 12 as a function of forecast horizon.

5.3. Exploration of Forecast Errors

The forecast skill of the network-based forecasts remains at a surprising +20% at time horizons through 2 h. This was unexpected because the finite domain of the network is usually transited by clouds in 10 to 20 min. To explain this finding, we revisited the underlying statistics of forecast skill. The root mean squared error can be written as

$$\text{RMSE} = \sqrt{\sigma_f^2 + \sigma_o^2 - 2\sigma_f\sigma_o\rho + \text{MBE}^2} \quad (11)$$

where σ_f is the forecast standard deviation, σ_o is the measurement standard deviation, ρ is the correlation coefficient, and MBE is the mean bias error (Taylor, 2001). When correlations and biases are small, the RMSE reduces to a sum in quadrature of the observation and measurement standard deviations. Under these conditions, a smoother forecast will have a lower RMSE, and thus a more positive forecast skill, than a more variable forecast. Of course, this does not mean that the smoother forecast is more skillful under most definitions of the word.

As an alternative means of understanding the relative merits of our forecast methods, we turned to Taylor diagrams (Taylor, 2001). The Taylor diagram in Fig. 13 shows the CRMSE, correlation coefficient, and standard deviations of clear-sky index forecasts for each forecast method. Here, we analyzed forecasts of clear-sky index instead of irradiance so all values are dimensionless. The solid contour lines are lines of constant CRMSE. We

Figure 10: RMSE of many types of forecasts averaged over 46 cloudy days. Clear sky refers to a forecast where one assumes the sky is always clear ($k_n^*(t) = 1$).

see that network forecasts have correlations greater than or approximately equal to spatially-averaged persistence but with higher standard deviation. This means network forecasts capture more variability. Network forecast standard deviation transitions from performing like clear-sky index persistence forecasts at short time horizons to approaching spatially-averaged persistence, analogous to the transitions for MAE and RMSE in Figs. 9 and 10. At roughly 30 min forecast horizons, network forecasts behave about the same as spatially-averaged persistence forecasts as we expect based on the method used and average cloud velocities. Hence, we say that our network forecasts are more useful than simple spatial averaging for forecast horizons less than 30 min. Regardless of their forecast skill metric scores, assessing the utility of network and spatial-average persistence forecasts past 30 min is challenging. We therefore suggest that researchers restrict their use of forecast skill to methods which have similar mean bias errors and standard deviations.

Figure 13 also shows how network and spatially-averaged persistence forecasts always have lower RMSE than clear-sky index persistence after a certain horizon. This is a result of the combination of lower standard deviation and higher correlation for the network and spatially-averaged persistence forecasts. This trend holds for even longer forecast horizons. Unfortunately, Eq. (11) does not simplify for the forecasts and data shown here so both correlation and standard deviation need to be considered to understand RMSE.

5.4. Limitations and Comparisons to Other Work

One limitation of the current network algorithm is that it does not account for multiple cloud layers. Satellite images from many of the studied days confirm that multiple cloud layers were moving in different directions. We also studied incorporating data from times in the past appropriately shifted by cloud motion vectors but found no noticeable improvement, likely due to this complex motion.

On a day with a single cloud layer coming from the southwest shown in Fig. 14, we see that a single upstream sensor greatly improves network forecasts at around the 7 min forecast horizon. This demonstrates that the network method can perform quite well if the velocity of the clouds is well defined and the sensors are appropriately located.

Another limitation is the size of the irradiance network. Depending on the wind motion vectors clouds can pass from the edge of the network to the center in 10 min. Since the boundary is set to the spatial average of sensors, network forecasts converge to spatially averaged persistence.

Still, our current method of network forecasting performs as well as or better than both clear-sky index and spatially-averaged persistence. Error statistics for network forecasts for cloudy days are presented in Table 1.

When we compared our current network method and high resolution data with the previous work of Lonij et al. (2013), we see that our new method performs favorably. Lonij et al. use a network of 80 rooftop PV systems in the Tucson area with 15 min averaged power data to make short-term forecasts of power. Their method uses a similar cloud translation method as this work, but wind vectors are obtained from NOAA forecasts, via optimization of the wind vector to minimize RMS forecast errors, or via a Kalman filter applied to optimized vectors. At 15 and 30 min forecast horizons, the best forecasts of Lonij et al. had skills of -8.0% and 2.4%, respectively, while our new method has skills of 17.7% and 21.2%. Even compared to the optimized “forecasts” (which were not true forecasts) with skills of 1.6% and 34.5% at 15 and 30 min, our new method performs well. We only used 3 months of data from our real-time network while Lonij et al. used one year of data.

Chu et al. (2015b) produced a cloud tracking forecast of PV power with an ANN applied to a deterministic forecast using a sky imager at a site near the Nevada/Arizona border. The initial deterministic forecast model does not perform well compared to persistence, with negative skills at 5, 10, and 15 min forecast horizons. However, the re-forecast using an ANN technique improves the result with skills of 15.1%, 21.8%, and 26.2% at forecast horizons of 5, 10, and 15 min respectively, which are

comparable to our technique. Similar optimization could be applied to our deterministic network forecasts to further improve skill. A Taylor diagram of both the initial deterministic forecast and ANN re-forecast would be useful as another method to assess the forecasts.

Compared to the regression methods in Yang et al. (2015), our forecasts perform comparably at the 5 min forecast horizon. Yang et al. used 1 s irradiance data from Oahu and applied the lasso and ordinary least squares regression methods to make very short term (< 5 min) forecasts. At shorter horizons, both methods can outperform the reference persistence forecast. Since our forecasts approach the clear-sky index persistence model, regression methods are likely a better choice if sub-five minute time horizons forecasts are needed, at least for the region studied here.

6. Conclusion

We presented a deterministic method to forecast irradiance that uses data from a network of irradiance sensors as the primary input. This method can combine the benefits of clear-sky index persistence and spatially-averaged persistence into one forecast. It outperforms a reference clear-sky index persistence model for 1 to 120 min forecast horizons. Much of this improvement is due to spatial averaging, which shows surprising utility for the region and time period studied. However, network forecasts still exhibit more variability than spatially-averaged persistence, thus we claim network forecasts are better at forecasts horizons less than 30 min. The results presented here used numerical

Table 1: Summary of error statistics for network forecasts for the 46 days with clouds. Error statistics were calculated for the entire dataset at once. Only forecasts and data with solar zenith angle less than 75° were used. The mean irradiance was $\bar{y} = 662 \text{ W/m}^2$ and the mean clear-sky index was $\bar{k} = 0.92$.

FH	rMAE (%)	MAE (W/m^2)	MBE (W/m^2)	rRMSE (%)	RMSE (W/m^2)	Avg. skill (%)
1 min	4.96	30.97	-1.44	11.90	82.55	22.96
3 min	7.51	48.13	-1.39	15.89	110.46	23.09
5 min	9.29	59.59	-3.91	18.67	127.06	19.65
10 min	11.39	71.38	-8.59	22.11	141.44	18.63
20 min	13.23	82.39	-10.46	24.03	152.84	18.66
30 min	13.95	86.57	-7.52	24.49	154.15	21.21
60 min	15.45	95.59	-6.65	26.59	160.72	21.00
120 min	17.02	106.51	-2.01	29.20	172.45	19.58

weather model winds at a single layer of the atmosphere to perform cloud advection, so complex cloud movement or incorrect cloud motion vectors likely limited the accuracy. The limited size and density of the network also limits the accuracy of network forecasts.

We showed that forecast skill can be a misleading metric, and we instead used a Taylor diagram to better understand the differences among forecast methods. This lead us to reinterpret our finding that network forecasts show significant skill to 2 h forecast horizons so now we make a more informed claim that network forecasts show meaningful skill out to 30 min forecast horizons. We encourage other authors to make use of Taylor diagrams when assessing the quality of forecasts.

While the method presented may have a limited useful maximum forecast horizon, the irradiance sensor network will be a valuable asset to make other types of forecasts. For instance, regression methods using a network can improve very short

time horizon forecasts (Yang et al., 2015). In the future, we could use the network of sensors to improve satellite image forecasts similar to Marquez et al. (2013) and to validate numerical weather model forecasts. We may also study how different interpolation methods affect the results of our network-based forecasting method in a detailed comparison.

Acknowledgments

This work was supported by Tucson Electric Power (TEP). We thank Mike Leuthold from the University of Arizona Department of Atmospheric Sciences for providing WRF data, and the University of Arizona Renewable Energy Network (UAREN) for institutional support. We also thank Technicians for Sustainability, a local solar PV installer, for providing data from rooftop PV systems in near real-time for forecasting and analysis. WFH thanks the Department of Energy (DOE) Office of Energy Efficiency and Renewable Energy (EERE) Postdoctoral Research Award for support.

Figure 11: Network RMSE vs clear-sky index persistence RMSE for all days and 5, 10, 15, and 30 min forecast horizons calculated for clear-sky indices. The slope of the best fit line is used to approximate the skill of the forecasts. Each point represents one day of data and forecasts. The forecast skill is positive for any point below the $y = x$ line. The plots illustrate that network forecasts may have negative skill for days that are nearly clear. Furthermore, they show a consistent positive skill for cloudier days with few outliers.

Figure 12: Average skill of our network forecasts for the study period of April, May, and June 2014.

Figure 13: Taylor diagram for clear-sky index persistence (red), spatially-averaged persistence (yellow), 5-min time-averaged persistence (green), and network (light blue) forecasts for 1 min (circle), 3 min (triangle), 7 min (diamond), 30 min (square), and 120 min (pentagon) forecast horizons. The black dashed line indicates the standard deviation of the data. Solid contours around the observations point are lines of constant CRMSE. Forecasts for clear-sky index were used so all quantities are dimensionless. At the 120 min forecast horizon, the spatially-averaged persistence and network points overlap. Network forecasts start with a standard deviation near that of the measurements, but this decreases at longer time horizons as the network forecast begins to resemble spatially-averaged persistence.

Figure 14: RMSE vs forecast horizon on May 19, 2014 for network forecasts made with all the sensors in the network (blue) and with one upwind sensor removed (red), along with a spatially-averaged persistence forecast (yellow). The dip at 7 min for the forecast using the full network illustrates that properly placed upstream sensors do improve forecasts over a simple spatial average.

References

- Achleitner, S., Kamthe, A., Liu, T., Cerpa, A. E., Apr. 2014. SIPS: Solar Irradiance Prediction System. IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks, 225–236. URL <http://dx.doi.org/10.1109/IPSIN.2014.6846755>
- Bilionis, I., Constantinescu, E. M., Anitescu, M., Dec. 2014. Data-driven model for solar irradiation based on satellite observations. *Solar Energy* 110, 22–38. URL <http://dx.doi.org/10.1016/j.solener.2014.09.009>
- Bosch, J., Zheng, Y., Kleissl, J., Jan. 2013. Deriving cloud velocity from an array of solar radiation measurements. *Solar Energy* 87, 196–203. URL <http://dx.doi.org/10.1016/j.solener.2012.10.020>
- Chu, Y., Pedro, H. T., Li, M., Coimbra, C. F., Apr. 2015a. Real-time forecasting of solar irradiance ramps with smart image processing. *Solar Energy* 114, 91–104. URL <http://dx.doi.org/10.1016/j.solener.2015.01.024>
- Chu, Y., Urquhart, B., Gohari, S. M., Pedro, H. T., Kleissl, J., Coimbra, C. F., Feb. 2015b. Short-term reforecasting of power output from a 48 MWe solar PV plant. *Solar Energy* 112, 68–77. URL <http://dx.doi.org/10.1016/j.solener.2014.11.017>
- Cormode, D., 2015. Large and Small Photovoltaic Powerplants. Phd dissertation, University of Arizona.
- Diagne, M., David, M., Boland, J., Schmutz, N., Lauret, P., Jul. 2014. Post-processing of solar irradiance forecasts from WRF model at Reunion Island. *Solar Energy* 105, 99–108. URL <http://dx.doi.org/10.1016/j.solener.2014.03.016>
- Elsinga, B., van Sark, W. G., 2014. Inter-system time lag due to clouds in an urban PV ensemble. In: 2014 IEEE 40th Photovoltaic Specialist Conference (PVSC). No. 4. pp. 0754–0758. URL <http://dx.doi.org/10.1109/PVSC.2014.6925029>
- Franke, R., 1982. Scattered data interpolation: tests of some methods. *Mathematics of Computation* 38 (157), 181–181. URL [http://dx.doi.org/10.1090/S0025-5718-1982-0637296-4](http://dx.doi.org/10.1090/amt-7-1693-2014)
- Fung, V., Bosch, J. L., Roberts, S. W., Kleissl, J., Jun. 2014. Cloud shadow speed sensor. *Atmospheric Measurement Techniques* 7 (6), 1693–1700. URL <http://dx.doi.org/10.5194/amt-7-1693-2014>
- Gueymard, C. A., Mar. 2008. REST2: High-performance solar radiation model for cloudless-sky irradiance, illuminance, and photosynthetically active radiation Validation with a benchmark dataset. *Solar Energy* 82 (3), 272–285. URL <http://dx.doi.org/10.1016/j.solener.2007.04.008>
- Hammer, A., Heinemann, D., Lorenz, E., Lückehe, B., Jul. 1999. Short-term forecasting of solar radiation: a statistical approach using satellite data. *Solar Energy* 67 (1-3), 139–150. URL [http://dx.doi.org/10.1016/S0038-092X\(00\)00038-4](http://dx.doi.org/10.1016/S0038-092X(00)00038-4)
- Hill, C. A., Such, M. C., Chen, D., Gonzalez, J., Grady, W. M., 2012. Battery energy storage for enabling integration of distributed solar power generation. *IEEE Transactions on Smart Grid* 3 (2), 850–857. URL <http://dx.doi.org/10.1109/TSG.2012.2190113>
- Ineichen, P., Perez, R., Sep. 2002. A new airmass independent formulation for the Linke turbidity coefficient. *Solar Energy* 73 (3), 151–157. URL [http://dx.doi.org/10.1016/S0038-092X\(02\)00045-2](http://dx.doi.org/10.1016/S0038-092X(02)00045-2)
- Inman, R. H., Pedro, H. T., Coimbra, C. F., Dec. 2013. Solar forecasting methods for renewable energy integration. *Progress in Energy and Combustion Science* 39 (6), 535–576. URL <http://dx.doi.org/10.1016/j.peccs.2013.06.002>
- Joskow, P. L., May 2011. Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies. *American Economic Review* 101 (3), 238–241. URL <http://dx.doi.org/10.1257/aer.101.3.238>
- Kleissl, J. (Ed.), 2013. *Solar Energy Forecasting and Resource Assessment*, 1st Edition. Elsevier, Oxford. URL <http://dx.doi.org/10.1016/B978-0-12-397177-7.18001-5>
- Lauret, P., Diagne, M., David, M., 2014. A Neural Network Post-processing Approach to Improving NWP Solar

Radiation Forecasts. Energy Procedia 57, 1044–1052.

URL <http://dx.doi.org/10.1016/j.egypro.2014.10.089>

Lave, M., Kleissl, J., May 2013. Cloud speed impact on solar variability scaling Application to the wavelet variability model. Solar Energy 91, 11–21.

URL <http://dx.doi.org/10.1016/j.solener.2013.01.023>

Leuthold, M., 2015. The University of Arizona Department of Atmospheric Sciences - Arizona Regional WRF Model Data.

URL <http://www.atmo.arizona.edu/index.php?section=weather&id=wrf>

Lipperheide, M., Bosch, J., Kleissl, J., Feb. 2015. Embedded nowcasting method using cloud speed persistence for a photovoltaic power plant. Solar Energy 112, 232–238.

URL <http://dx.doi.org/10.1016/j.solener.2014.11.013>

Lonij, V. P., Brooks, A. E., Cronin, A. D., Leuthold, M., Koch, K., Nov. 2013. Intra-hour forecasts of solar power production using measurements from a network of irradiance sensors. Solar Energy 97, 58–66.

URL <http://dx.doi.org/10.1016/j.solener.2013.08.002>

Lorenzo, A. T., Cronin, A. D., Holmgren, W. F., Aug. 2015. Irradiance monitoring network data and wind motion vectors.

URL <http://zenodo.org/record/29070>

Lorenzo, A. T., Holmgren, W. F., Leuthold, M., Kim, C. K., Cronin, A. D., Betterton, E. a., 2014. Short-term PV power forecasts based on a real-time irradiance monitoring network. 2014 IEEE 40th Photovoltaic Specialist Conference (PVSC), 0075–0079.

URL <http://dx.doi.org/10.1109/PVSC.2014.6925212>

Marquez, R., Coimbra, C. F. M., Oct. 2012. Proposed Metric for Evaluation of Solar Forecasting Models. Journal of Solar Energy Engineering 135 (1), 011016.

URL <http://dx.doi.org/10.1115/1.4007496>

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

URL <http://dx.doi.org/10.1016/j.solener.2013.02.023>

023

Mathiesen, P., Kleissl, J., May 2011. Evaluation of numerical weather prediction for intra-day solar forecasting in the continental United States. Solar Energy 85 (5), 967–977.

URL <http://dx.doi.org/10.1016/j.solener.2011.02.013>

Nuss, W. A., Titley, D. W., Jul. 1994. Use of Multiquadric Interpolation for Meteorological Objective Analysis. Monthly Weather Review 122 (7), 1611–1631.

URL [http://dx.doi.org/10.1175/1520-0493\(1994\)122<1611:UOMIFM>2.0.CO;2](http://dx.doi.org/10.1175/1520-0493(1994)122<1611:UOMIFM>2.0.CO;2)

Perez, R., Kivalov, S., Schlemmer, J., Hemker, K., Renné, D., Hoff, T. E., Dec. 2010. Validation of short and medium term operational solar radiation forecasts in the US. Solar Energy 84 (12), 2161–2172.

URL <http://dx.doi.org/10.1016/j.solener.2010.08.014>

Perez, R., Lorenz, E., Pelland, S., Beauharnois, M., Van Knowe, G., Hemker, K., Heinemann, D., Remund, J., Müller, S. C., Traunmüller, W., Steinmauer, G., Pozo, D., Ruiz-Arias, J. a., Lara-Fanego, V., Ramirez-Santigosa, L., Gaston-Romero, M., Pomares, L. M., Aug. 2013. Comparison of numerical weather prediction solar irradiance forecasts in the US, Canada and Europe. Solar Energy 94, 305–326.

URL <http://dx.doi.org/10.1016/j.solener.2013.05.005>

Sirayanone, S., 1988. Comparative studies of kriging, multiquadric-biharmonic and other methods for solving resources problems. Dissertation, Iowa State University.

Taylor, K. E., 2001. Summarizing multiple aspects of model performance in a single diagram. Journal of Geophysical Research: Atmospheres 106 (D7), 7183–7192.

URL <http://dx.doi.org/10.1029/2000JD900719>

Urquhart, B., Ghonima, M., Nguyen, D. A., Kurtz, B., Chow, C. W., Kleissl, J., 2013. Sky-Imaging Systems for Short-Term Forecasting. In: Solar Energy Forecasting and Resource Assessment. Elsevier, Ch. 9, pp. 195–232.

URL <http://dx.doi.org/10.1016/B978-0-12-397177-7.00009-7>

Webster, R., Oliver, M. A., 1993. How large a sample is needed to estimate the regional variogram adequately? In:

- 1
- 2
- 3
- 4 765 Soares, A. (Ed.), Geostatistics Tróia 92 SE - 14. Vol. 5 of
- 5 Quantitative Geology and Geostatistics. Springer Nether-
- 6 lands, pp. 155–166.
- 7 URL http://dx.doi.org/10.1007/978-94-011-1739-5_14
- 8 14
- 9
- 10 770 Wilcox, S., Andreas, A., Nov. 2010. Solar Resource & Meteo-
- 11 rological Assessment Project (SOLRMAP): Observed At-
- 12 mospheric and Solar Information System (OASIS); Tuc-
- 13 son, Arizona (Data).
- 14 URL <http://dx.doi.org/10.7799/1052226>
- 15
- 16 775 Yang, D., Ye, Z., Lim, L. H. I., Dong, Z., 2015. Very short
- 17 term irradiance forecasting using the lasso. Solar Energy
- 18 114, 314–326.
- 19 URL <http://dx.doi.org/10.1016/j.solener.2015.01.016>
- 20
- 21
- 22 780 Yang, H., Kurtz, B., Nguyen, D., Urquhart, B., Chow,
- 23 C. W., Ghonima, M., Kleissl, J., May 2014. Solar irra-
- 24 diance forecasting using a ground-based sky imager de-
- 25 veloped at UC San Diego. Solar Energy 103, 502–524.
- 26 URL <http://dx.doi.org/10.1016/j.solener.2014.02.044>
- 27
- 28 785 044
- 29
- 30 Zhang, J., Florita, A., Hodge, B.-M., Lu, S., Hamann, H. F.,
- 31 Banunarayanan, V., Brockway, A. M., Jan. 2015. A suite
- 32 of metrics for assessing the performance of solar power
- 33 forecasting. Solar Energy 111, 157–175.
- 34
- 35 790 URL <http://dx.doi.org/10.1016/j.solener.2014.10.016>
- 36
- 37
- 38
- 39
- 40
- 41
- 42
- 43
- 44
- 45
- 46
- 47
- 48
- 49
- 50
- 51
- 52
- 53
- 54
- 55
- 56
- 57
- 58
- 59
- 60
- 61
- 62
- 63
- 64
- 65