

Short-term solar irradiance forecasting under data transmission constraints

Joshua E. Hammond^a, Ricardo A. Lara Orozco^b, Michael Baldea^{a,d,*}, Brian A. Korgel^{a,c,e}

^a McKetta Department of Chemical Engineering, The University of Texas at Austin, 200 East Dean Keeton St., Stop C0400, Austin, 78712, TX, United States

^b Hildebrand Department of Petroleum and Geosystems Engineering, The University of Texas at Austin, 200 E. Dean Keeton St., Stop C0300, Austin, 78712, TX, United States

^c Energy Institute, The University of Texas at Austin, 2304 Whitis Ave., Stop C2400, Austin, 78712, TX, United States

^d Institute for Computational Engineering and Sciences, The University of Texas at Austin, 201 E. 24th St., POB 4.102, Stop C0200, Austin, 78712, TX, United States

^e Texas Materials Institute, The University of Texas at Austin, 204 E. Dean Keeton St., Stop C2201, Austin, 78712, TX, United States

ARTICLE INFO

Dataset link: <https://github.com/joshuaeh/TabularSolarForecast>

Keywords:

Feature engineering
Forecasting
Machine learning
Renewable energy
Solar irradiance

ABSTRACT

We report a data-parsimonious machine learning model for short-term forecasting of solar irradiance. The model follows the convolutional neural network – long-short term memory architecture. Its inputs include sky camera images that are reduced to scalar features to meet data transmission constraints. The model focuses on predicting the deviation of irradiance from the persistence of cloudiness (POC) model. Inspired by control theory, a noise signal input is used to capture the presence of unknown and/or unmeasured input variables and is shown to improve model predictions, often considerably. Five years of data from the NREL Solar Radiation Research Laboratory were used to create three rolling train-validate sets and determine the best representations for time, the optimal span of input measurements, and the most impactful model input data (features). For the chosen validation data, the model achieves a mean absolute error of 74.29 W/m² over a time horizon of up to two hours, compared to a baseline 134.35 W/m² using the POC model.

1. Introduction

The United States has set an aggressive target to achieve a power grid with net-zero greenhouse gas emissions by 2035 [1–3]. This will require a major shift in the power generation mix from fossil fuels, to include significantly more renewable sources, such as wind and solar [3–5]. Since wind and solar power generation are intermittent and largely non-dispatchable, it will become increasingly important to anticipate fluctuations in renewable power generation to ensure the stability of the grid. Accurate near-term forecasts of solar photovoltaic (PV) power generation will be especially important.

The economic impact of PV forecasts is considerable [6]. Accurate near-term forecasts enhance grid stability by anticipating ramping events [7–10]. In addition, forecasts may be used to define targets for demand response, and stabilize grid dynamics [11,12], advise real-time market price predictions [13], and inform ancillary service dispatch [14,15].

Short-term forecasts of PV generation often rely on large volumes of diverse, high-dimensional data including local meteorological measurements, numerical weather prediction [16], and satellite images [17]. Sky cameras are becoming an increasingly widespread method for gathering local information at generation sites [18–20]. Existing irradiance

forecasting work integrates these heterogeneous data sources to improve forecast accuracy amid the largely stochastic weather processes that govern irradiance.

Local measurements are usually transmitted and processed at a centralized facility or data center to produce forecasts because of the computational power required. For example, in the Electric Reliability Council of Texas (ERCOT), local meteorological data are collected at all grid-connected PV generation sites and transmitted to inform operations and produce generation forecasts in a centralized fashion [16]. However, solar PV facilities are typically placed in remote locations, and transmitting the required data (particularly, sky camera images) from the generation sites to centralized computing facilities can be difficult due to data infrastructure costs and data transmission capacity limitations.

It is precisely this problem of forecasting under data transmission constraints that we address in this work. We introduce a data-parsimonious machine learning model to forecast solar irradiance up to two hours ahead. The model uses a set of meteorological station data and *scalar* data extracted from sky camera images as inputs to minimize the data transmission required. To our knowledge, this

* Corresponding author at: McKetta Department of Chemical Engineering, The University of Texas at Austin, 200 East Dean Keeton St., Stop C0400, Austin, 78712, TX, United States.

E-mail addresses: mbaldea@che.utexas.edu (M. Baldea), korgel@che.utexas.edu (B.A. Korgel).

<https://doi.org/10.1016/j.renene.2024.121058>

Received 28 March 2024; Received in revised form 28 June 2024; Accepted 21 July 2024

Available online 23 July 2024

0960-1481/© 2024 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

Nomenclature

ASI	All Sky Imager
BMS	Baseline Measurement System
CNN	Convolutional Neural Network
CS Dev.	Clear Sky Deviation
ERCOT	Electric Reliability Council of Texas
FSS	Forecast Skill Score
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
nMAP	Normalized Mean Absolute Percentage Error
NREL	National Renewable Energy Laboratory
NWP	Numerical Weather Prediction
POC	Persistence of Cloudiness
PV	Photovoltaic
RMSE	Root Mean Square Error
SRRL	Solar Radiation Research Laboratory

Time Representations

$\angle$	Sine and Cosine Transformation
TM	Time Milestones
ToD	Time of Day
ToY	Time of Year

Irradiance Representations

ΔCSI	Change in Clear Sky Index
ΔGHI	Change in Irradiance
$\hat{y}_i$	A predicted value
y_i	A true value
CS Dev.	Clear Sky Deviation
CSI_t	Clear Sky Index at time t
DHI	Diffuse Horizontal Irradiance
DNI	Direct Normal Irradiance
GHI	Global Horizontal Irradiance
GHI_0	Irradiance at time of forecast
GHI_t	Irradiance at time t
$GHI_{CS,t}$	Ideal Clear Sky Irradiance at time t

is the first work that considers the limitations of data transmission during model development. In this work, feature importance studies are used to identify the most impactful features from the available set. Inspired by control theory, a noise signal (along with the corresponding noise model) is introduced to capture unmeasured and/or unknown variables and disturbances that impact irradiance. The model output (predicted variable) is the residual (difference) between the persistence of cloudiness (POC) prediction and true irradiance over the forecast horizon. While the final model is less accurate than recently published sky camera-based forecasting models, it requires orders of magnitude less data transmission and significantly outperforms reference forecasts.

The specific contributions of this work are:

- A parsimonious machine learning model for near-term forecasting of solar irradiance using scalar features. This model is designed to minimize data transmission requirements from remote solar PV generation sites and requires significantly less data bandwidth than existing models.
- An empirical comparison of model performance using different irradiance representations. Predicting deviation from the POC

model eliminates the need to capture known long-term dynamics, and is shown to improve forecast accuracy (compared to including persistence information as a model input).

- A noise model is used for the first time in this application area to the best of our knowledge, to account for unmeasured and/or unknown variables/disturbances, and is shown to further improve forecasting accuracy.
- New empirical insights on feature importance and the effect of input sequence length on model performance are drawn from a large-scale data set.

2. Background

The term “solar forecasting” describes the prediction of PV power generation, as well as the prediction of irradiance, given that PV generation is a function of global horizontal irradiance (GHI, W/m^2) and panel temperature [21]. While cloud cover and position have the largest impact on irradiance [22–24] — sometimes accounting for a change in irradiance of over 80% in a minute — other variables such as dispersed particles and wind [25] are linked to irradiance as well. As shown in Eq. (1), GHI is composed of both direct and diffuse sunlight, where DNI is the direct normal irradiance, α is the solar zenith angle, and DHI is the diffuse horizontal irradiance. Clear sky models calculate GHI in the absence of weather effects using time and global position [21,26–28]. The ratio of measured irradiance (GHI_t) to ideal clear sky irradiance ($GHI_{CS,t}$) at a given time instant t is often referred to as clear sky index (CSI) and is shown in Eq. (2).

$$GHI = \cos(\alpha)DNI + DHI \quad (1)$$

$$CSI_t = \frac{GHI_t}{GHI_{CS,t}} \quad (2)$$

Persistence models are often used to evaluate the effectiveness of forecasting algorithms. The persistence of cloudiness (POC) model assumes that the clear sky index at the time of forecast will remain constant throughout the forecast horizon. Eq. (3) shows an example POC forecast at 10-min intervals up to a maximum forecast horizon of 120 min.

$$GHI_t = CSI_0 \cdot GHI_{CS,t}, t \in \{10, 20, \dots, 110, 120\} \quad (3)$$

The POC forecast does not take into account fluctuations in weather conditions, but does provide a baseline for near-term solar irradiance forecasting. More sophisticated forecast models should therefore achieve higher accuracy compared to the POC model by anticipating changes in weather conditions.

It is common for these latter models to predict the *difference* between the time-varying irradiance and the baseline. This has the effect of de-trending the data [29,30]. For example, [31,32] use autoregressive models to predict irradiance relative to a clear sky model (rather than using the POC as a baseline). We note that other de-trending methods, such as wavelet decomposition [33] and spectral decomposition [34] have been recently proposed for irradiance forecasting.

The aforementioned works are restricted to using *scalar* data as model inputs, and the number of inputs is typically small in practical settings.

Nevertheless, there is a wealth of data relevant to irradiance forecasting. These include local meteorological measurements (cloud height, wind speed and direction, temperature, humidity, and air pressure), sky camera images, Numerical Weather Prediction (NWP) results, and satellite imaging. Some of these data are in non-scalar formats (e.g., camera images) and thus do not lend themselves naturally to use in time-series models.

This fact has motivated the use of machine learning (ML) approaches for irradiance forecasting. Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks are common

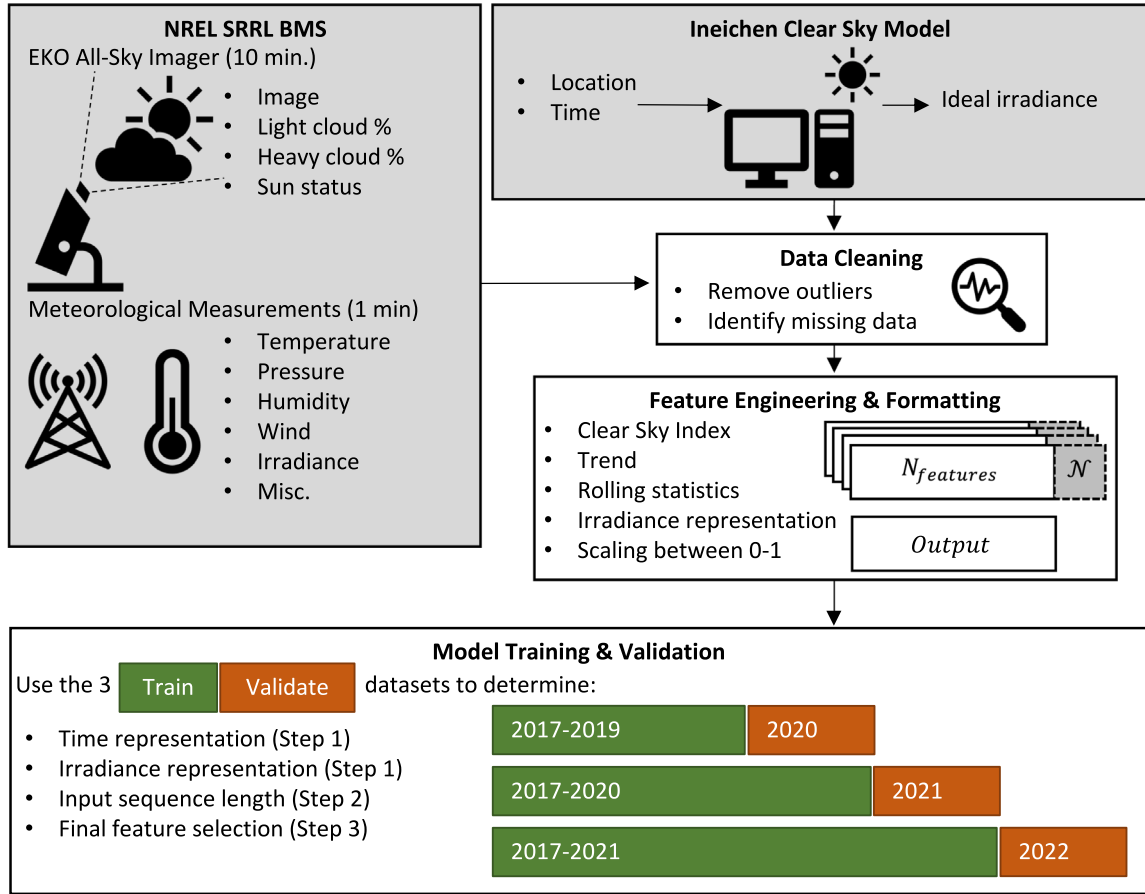


Fig. 1. Data collection, curation and model training workflow.

deep learning architectures used to extract patterns from data and predict sequences, respectively [35].

ML models broadly fall into two categories:

- End-to-end models utilize available data as inputs with little or no pre-processing or interpretation. For example, Peng et al. [36] reported the use of multiple sky cameras to track cloud positions in the sky, and a stereoscopic arrangement was reported by Nouri et al. [37] for estimating cloud height. Song et al. [38] determined current irradiance directly from sky camera images, while Lin et al. [39] used deep learning models to forecast irradiance from similar data. Some end-to-end models are used for multiple predictions, e.g., both cloud patterns [40–42] and irradiance [19,43,44]. Liu et al. [45] use a deep learning model to determine cloud movement vectors from image sequences as well as forecast irradiance. These models typically use heterogeneous data formats as inputs. It is worth noting the work of Ogliari et al. [46], who performed feature fusion by leveraging unused image pixels to embed other sensor data before feeding sky-camera images to a deep CNN model.
- Physically-motivated models take a physics-based lens to the data, generating additional insights prior to building the irradiance forecasting model itself. Ideally, physics-based models should reduce the parameter search space during model training and show better generalization performance. Physics-based approaches include detection of cloud position, cloud motion, and identifying position of the Sun, which can be approached via color analysis [47,48], pixel clustering [49,50], and convolution methods [51]. Recent work by Paletta et al. [17] demonstrated that polar transformations centered on the sun can extract cloud

movement toward the sun. Fabel et al. [52] use features extracted from a sky camera such as cloud coverage, height, and type as well as parametric relationships such as solar position to predict future irradiance.

We emphasize here that these are *broad* categories and modeling can take inspiration from both. For example, Le Guen and Thome [20] use physics-based fluid-flow equations to predict future irradiance using sky camera images as inputs in an end-to-end structure.

We also note that the works referenced here use sophisticated, computationally-intensive and data-intensive architectures to perform the irradiance forecasting task. Local data collected from — typically remote — solar PV generation sites are transmitted over some communication channel to a centralized location for processing and forecasting. This implicitly assumes that there are no data transmission constraints. In practice, however, the communication channel may have severely limited bandwidth, and may not be able to accomplish the timely transmission of all the data—particularly sky camera images which are typically stored in large files. This emphasizes the need for developing *data-parsimonious* models, capable of performing the forecasting task based on a *limited* number of inputs that are readily obtainable over existing bandwidth-restricted communication channels.

3. Methods

Motivated by bandwidth-constrained applications, we present a CNN-LSTM model which uses *scalar* features extracted from sky camera images to forecast irradiance. In Section 3.1 we describe the data used and pre-processing steps. Section 3.2 defines the model architecture including the noise model, and Section 3.3 contains the training and evaluation procedures used.




Fig. 2. Irradiance measurements and a subset of sky-camera images from May 18, 2022 at the NREL SRRL BMS. Figure created by the authors based on data collected from Stoffel and Andreas [53]. Letters denote the moments that the images were collected.

3.1. Data

We begin with measurements spanning across five years (2017–2022) from the National Renewable Energy Laboratory (NREL) Solar Radiation Research Laboratory (SRRL) Baseline Measurement System (BMS) [53] in Golden, Colorado, USA. Fig. 1 illustrates the flow of the data from primary data sources to the final model training and validation process. The NREL SRRL BMS provides 131 parameters measured each minute from an onsite meteorological station including temperature, barometric pressure, wind speed, and irradiance. The BMS also includes images taken by an Eko All Sky Imager (ASI-16) at 10-min intervals. The scalar features extracted from those images include light and heavy cloud coverage percentages and whether the sun is covered by clouds [47,54]. Fig. 2 shows an example of measured irradiance values compared to clear sky irradiance values, as well as a selection of sky images at the time of measurement.

BMS data are augmented with ideal clear sky irradiance data [27, 28,55] using the Ineichen Clear Sky Model [26] provided by the pvlib package [55]. During the data cleaning process, data were clipped to physically valid values (such as imposing that irradiance values be non-negative), and outliers were removed if they had a z-score of over five. The feature engineering step included calculating values we hypothesized would be useful for the model to predict future irradiance. These calculated features include clear sky index and other irradiance and time representations, as well as lagged statistics for recent irradiance measurements. The dataset was also resampled to match the 10-min image frequency. Finally, the data were scaled to lie between 0 and 1 using the minimum and maximum values of each feature in the cleaned training set and formatted into input–output pairs. The inputs are two-dimensional matrices with rows representing discrete time instants 10 min apart and each column representing a feature. The outputs are the predicted irradiance values over the two-hour forecast horizon (recall that these are residuals (i.e., differences) between POC model predictions and the measured irradiance).

A full list of the scalar features extracted from sky camera images as well as a comprehensive list of all features examined in this work is provided in Appendix A. We note that the scalar features of the sky

camera images are provided by the camera itself and do not require any additional local processing or computing power. While we focus on predicting irradiance at ten minute intervals up to two hours from the time of forecast, we believe that the model can be easily adapted to other forecast frequencies and prediction horizons.

3.2. Model architecture and noise model

The model architecture is a CNN-LSTM structure as shown in Fig. 3. A dropout layer ensures that the model learns sparse and generalizable feature representations and helps prevent overfitting. Next, CNNs efficiently extract intermediate features from input data using convolutional filters learned during model training. The LSTM layer identifies sequential patterns within the training data and is used to identify temporal patterns within the data. The final dense layers perform a bottleneck operation and transform the intermediate representations to the final irradiance sequence prediction over the time horizon considered. We use two dense layers to ensure that the model has sufficient capacity to encode the predicted irradiance.

Similar architectures have been successfully used for irradiance prediction by e.g., Gao and Liu [56] and Paletta et al. [41]. Nevertheless, our proposed architecture include a novel feature, that, to our knowledge, has not been employed in the field of solar irradiance forecasting. Inspired by forecasting mechanisms used in the literature on model predictive control [57], a noise signal is included as an additional input. Noise models account for the unmeasured input variables/disturbances that affect a dynamical system, which are assumed to be random variables. Their statistical properties, as well as the gain (i.e. the magnitude of the impact in the model) assigned to each noise input, are specified based on knowledge of the system or are identified from the data. The adoption of a noise model in this work is supported by the fact that it is natural to expect that the available features do not fully explain or predict the variable of interest (i.e., irradiance). Equivalently, we hypothesize that changes in irradiance are subject to random, unmeasured inputs. In practice, it is typical to use a Gaussian noise signal with zero mean unit variance. In our model, the noise model is implemented by augmenting the measured data with an additional column of normally




Fig. 3. CNN-LSTM model structure with an optional noise model within the dashed box. The noise model features a random Gaussian variable and a trained response to account for random, unknown disturbances. Dimensions are shown in parentheses.

distributed random numbers and allowing the CNN-LSTM to learn the gain — the effect that random disturbances have on the output — rather than specifying it.

3.3. Training procedure

The model was trained using an expanding train-validate data set with three successive steps to refine the initial model. The five years of data are split up so that one full year is always used as the validation dataset in order to ensure that all phenomena within a year are equally represented. The training datasets expand at each step, rather than a random split so that the validation datasets always contain data that reflect days that are temporally located after the training data and realistically represent generalization error. All training data sets begin on September 27, 2017, and validation datasets begin the day after train datasets end.

Since repeated tests on the same data increase the likelihood of false positive results [58], we sequentially investigate key model parameters

within each dataset and implement the results during the next step.¹ The first step evaluated different representations of time and irradiance. The second step studied the effect of changing the time span of input measurements on validation error. The third step used the results from the previous two steps and implemented permutation feature importance [59] to identify the most impactful features for predictive performance. The training and validation data splits for each step are shown in Table 1. Each of these steps are described in more detail in the following subsections.

All training took place on a single NVIDIA RTX 3090 Ti GPU with 24 GB of memory and batch size of 4000. The model was trained using the Adam optimizer [60] with an initial learning rate of 0.01 with cosine decay [61] over 500 epochs repeated 4 times. Early stopping was permitted if the validation loss did not decrease for 1000 epochs.

3.3.1. Step 1 — Time and irradiance representations

Time representations include a floating point number between 0 and 1 representing the proportion of time that has passed since the start of a day or year. These are shown as the Time of Day (ToD) and Time of Year (ToY) representations in Eqs. (4) and (5). Time of day may also be represented relative to key moments in the day such as sunrise, solar noon, and sunset as shown in Eq. (6). We refer to a vector containing these three values as Time Milestones (TM). Finally, we include a trigonometric transform of ToD, ToY, and TM as is common for cyclic variables [62]. We denote these representations as $\angle$ ToY, $\angle$ TM, and $\angle$ ToD and their respective equations are shown in Eqs. (7), (8), and (9).

$$\text{ToD} = (\text{Hours} + \text{Minutes}/60 + \text{Seconds}/3600)/24 \quad (4)$$

$$\text{ToY} = (\text{Day of Year} + \text{ToD})/365 \quad (5)$$

$$\text{TM} = \begin{cases} \text{ToD}(\text{Time} - \text{Sunrise Time}), \\ \text{ToD}(\text{Time} - \text{Solar Noon Time}), \\ \text{ToD}(\text{Time} - \text{Sunset Time}) \end{cases} \quad (6)$$

$$\angle\text{ToD} = \{\sin(\text{ToD}), \cos(\text{ToD})\} \quad (7)$$

$$\angle\text{ToY} = \{\sin(\text{ToY}), \cos(\text{ToY})\} \quad (8)$$

$$\angle\text{TM} = \{\sin(\text{TM}), \cos(\text{TM})\} \quad (9)$$

Multiple representations of future irradiance were tested as target variables to determine which produced the most accurate forecast. These representations include GHI_t , CSI , CS Dev. (the difference between GHI_t and GHI_{CS} as shown in Eq. (10)), as well as two forms that are relative to the conditions at the time of forecast: change in irradiance (ΔGHI), and change in clear sky index (ΔCSI). These relative representations are shown in Eqs. (11) and (12) respectively. Differencing and autoregressive approaches similar to these relative representations of irradiance have been shown to improve the predictive ability of statistical models [29,30].

$$\text{CS Dev.} = GHI_{CS,t} - GHI_t \quad (10)$$

$$\Delta GHI = GHI_t - GHI_0 \quad (11)$$

$$\Delta CSI = CSI_t - CSI_0 \quad (12)$$

3.3.2. Step 2 — Input sequence length

LSTM models use sequential inputs and determine temporal patterns in those inputs to predict future values. During Step 2, we investigate how the number of input measurements affects model performance.

¹ If 20 tests were performed on the same dataset, with a p -value of 0.05, we would expect one false positive to occur. This is a problem known as multiple testing.

Table 1

Model training and tuning used three successive steps to refine the initial model with an input block of 131 features and 13 temporal measurements. Count refers to the number of complete input and output data blocks within each date range.

Step	Last training date	Count	Last validation date	Count	Description	Approximate split
1	Sep. 26, 2019	21,343	Sep. 26, 2020	12,509	Time and Irradiance	63:37
2	Sep. 26, 2020	33,852	Sep. 26, 2021	13,256	Input time horizon	72:28
3	Sep. 26, 2021	47,108	Sep. 26, 2022	9199	Feature importance	84:16




Fig. 4. Validation Mean Absolute Error (W/m²) across all predicted points ($t + 10$ to $t + 120$) for each model trained during Step 1 using year 3 as the validation set. In this graphic, each row represents a different combination for time representation and each column represents a method for irradiance representation.

3.3.3. Step 3 — Feature importance and noise model evaluation

To determine the most important features for the model, we used permutation feature importance [59] which measures the change in model performance when a feature is randomly permuted or corrupted. The importance of the feature is then evaluated by the increase in error. There are many methods for interpreting the subsequent results and determining which parameters are to be included or not. Often, this is done with a fixed value threshold for minimum importance. In contrast, this work seeks to minimize the number of features as well as prediction error. Consequently, we hypothesize that most of the features that only have marginal importance could be dropped and replaced with the noise model. We test this by dropping all but the most important parameters, then refitting the model and comparing model performance. To determine the effect of the noise model, we compare model performance across all three training-validation dataset splits with and without the noise model.

3.3.4. Evaluation metrics

During the training process, we use Mean Absolute Error (MAE) as the objective function to be minimized, shown in Eq. (13). Here, y_i represents the true value and $\hat{y}_i$ represents the predicted value. For an equitable comparison, we convert all predicted and true irradiance values back to GHI before comparing statistics. To compare to other models in the literature we also calculate Root-Mean Square Error (RMSE), normalized Mean Absolute Percentage error (nMAP), as well as Forecast Skill Score (FSS) which measures performance relative to the POC model. These are defined as follows in Eqs. (14), (15) and (16) respectively. We choose to train on MAE rather than RMSE to minimize

the effect of outliers in the training data. Final model performance across all of these statistics is reported in Table A.9.

$$MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \quad (13)$$

$$RMSE = \frac{1}{N} \sum_{i=1}^N \sqrt{(y_i - \hat{y}_i)^2} \quad (14)$$

$$nMAP = \frac{1}{N} \sum_{i=1}^N \frac{|y_i - \hat{y}_i|}{\frac{1}{N} \sum_{i=1}^N y_i} \times 100 \quad (15)$$

$$FSS = 1 - \frac{MAE}{MAE_{POC}} \quad (16)$$

4. Results

4.1. Step 1 — Time and irradiance representations

Fig. 4 shows the MAE of all predictions in validation set 1 over a 2-hour horizon. Each row represents a different representation of time and each column a different irradiance representation. The ToD and ToY representation of time and the ΔCSI representation of future irradiance achieve the lowest MAE across all predicted values in Step 1. Other time representations performed similarly, however, ToD and ToY allow the model to capture daily and yearly cyclical patterns without increasing the number of variables and increasing the likelihood of overfitting.

In conjunction with ToD and ToY as time representations, ΔCSI achieves the lowest error among the models tested in Step 1. All models trained with target variables that incorporate information from




Fig. 5. Probability density functions of the future irradiance representations. Note that the relative irradiance representations (ΔCSI and ΔGHI) are much more tightly distributed around the center, likely facilitating easier recognition of patterns which deviate from the present conditions.

the clear sky model outperform those that use GHI. Incorporating the known patterns of the Sun using the clear sky model allows the model to explicitly isolate the unknown short-term dynamics, such as increasing cloud cover, and better predict irradiance without the effect of confounding variables. Both relative measures of irradiance appear to be more tightly distributed around the center, as shown in Fig. 5. This likely facilitates easier recognition of patterns which deviate from the present conditions. If conditions typically persist, events with a high deviation from the present conditions are further from the typical values and may be better distinguished during the training process.

Fig. 6 shows the FSS of the best model for each future irradiance representation and demonstrates that models trained with relative representations of future irradiance are better able to anticipate changes in irradiance than models trained with other forms of irradiance. Of the two relative irradiance forms, Fig. 7 suggests that ΔCSI is likely a better target variable since the clear sky index is more strongly autocorrelated (as defined in Eq. (17)) at later lags than GHI or CS Dev.

$$r_s = \frac{\text{cov}(R(X), R(Y))}{\sigma_{R(X)} \sigma_{R(Y)}} \quad (17)$$

4.2. Step 2 — Input sequence length

Thirteen models were trained with input sequence lengths increasing from a single measurement to two hours of measurements captured in a 13-element sequence. The MAE as a function of time from the forecast is shown in Fig. 8 for all 13 models, with a darker hue indicating a longer input data sequence. The overall MAE across the two-hour prediction horizon is shown in Table 2 for both the training and validation datasets. The model with an input sequence length of three measurements achieves the lowest training error, while the model with a single input measurement achieves the lowest validation error.

Surprisingly, models trained with fewer input data measurements have lower error than models that include more lagged information. However, Feng et al. [19] report a similar finding when forecasting irradiance at the same NREL SRRL BMS site and sky camera with an optimal input sequence length of two, which was the shortest sequence tested. The increase in error when moving from training to the validation set for the model trained with three input measurements suggests that the model is overfitting to the training data. The model training process can be approached as a combinatorial optimization problem.

Table 2

The overall MAE across the two-hour prediction horizon for each model trained during Step 2. The model trained with a single input measurement achieves the lowest validation error which is a better measure of true performance than training error.

Input sequence length	Train MAE (W/m ²)	Validation MAE (W/m ²)
1	81.29	76.53
2	79.25	79.06
3	75.96	81.78
4	81.77	80.86
5	85.57	82.02
6	83.03	82.84
7	82.81	84.64
8	86.68	84.49
9	90.23	85.00
10	88.95	84.98
11	90.70	86.45
12	85.85	87.38
13	97.26	86.97

In a predicament known as the “curse of dimensionality” [63], data-driven models require more and more data when increased input data and model parameters are included due to the search-space of solution combinations increasing. By providing less input data, the stochastic training process is less likely to result in overfitting or become stuck in local minima or saddle points.

Intuitively, these results indicate that future irradiance is more strongly associated with current weather than past measurements. Fig. 7 supports this conclusion, with autocorrelation consistently decreasing as the number of lags increase.

Because of the validation error and the minimal dimensionality, Step 3 focused on modeling with input sequence of one measurement — an input temporal horizon of 0 min.

4.3. Step 3 — Feature importance and noise model

Fig. 9 shows the results from Step 3, a feature importance test [59], which was used to determine which of the 168 features were most meaningful for the accuracy of GHI predictions. The test systematically corrupts one feature of the input data at a time, and observes the impact on model performance. Large increases in error indicate that a particular feature is valuable to the forecast accuracy while negligible




Fig. 6. FSS for a model trained on each future irradiance representation.




Fig. 7. Autocorrelation for three representations of future irradiance beginning 2 h or 12 lags before a prediction and extending over a two hour prediction horizon.

changes in error indicate that the model may not in fact rely on that feature.

The ten most important features identified in Step 3 are shown in Table 3 with the associated increase in MAE when they are corrupted. Rather than establish an arbitrary threshold for feature importance, we seek to minimize the number of features while maintaining a low prediction error. Since all but the top 10 features have an importance of less than 3, we retrained the model using only these 10 most important features, hypothesizing that the marginal importance of each remaining feature would be low and would be captured by the noise model.

Table 4 shows the validation error across all three train-validation data set splits both with and without the noise model. The architecture including the noise signal performs consistently better than the case without the noise model. Table 5 shows the results from a paired t-test comparing the model prediction error with and without the noise model to the null hypothesis that the noise model does not decrease error. Results are generally inconclusive for models with all features,

Table 3

Ten most important features of 131 total features as shown by the increase in predicted irradiance MAE as each feature is corrupted. Total cloud cover and GHI clear sky index contribute to prediction accuracy significantly more than other features.

Feature	Feature importance (MAE)
CDOC total cloud cover	37.76
CSI GHI	14.22
DNI _{t-4}	8.08
940 nm Aerosols	6.00
Mean CSI DNI Deviation	5.62
675 nm Aerosols	5.20
CSI DNI	4.88
Solar Elevation	4.73
CSI DNI Deviation	4.08
DNI _{t-9}	3.06

however the noise model consistently reduces the error of the model when provided with a reduced set of features. More importantly, using




Fig. 8. Mean Absolute Error (W/m^2) across all predicted points ($t+10$ to $t+120$) of models trained during Step 2 with increasing input time horizons. Each of these models used Δ CSI as the irradiance representation and ToD, ToY as the time representation based on the results in Step 1.




Fig. 9. Feature Importance distribution for the features studied.

Table 4

Validation error across all 3 steps of model refinement. Bold numbers indicate the best (or nearly-best) errors for each step. Results with and without the noise model as well as with all features and the top 10 features are shown. Results indicate that the noise model eliminates the drop in performance when the model is trained with fewer features.

Noise model	Features	Validation error (W/m^2)		
		Step 1	Step 2	Step 3
Not included	All	89.35	76.85	75.20
	Top 10	94.61	83.93	77.36
Included	All	91.07	76.01	74.95
	Top 10	88.96	76.01	74.29

the noise signal input eliminates the drop in performance when the model is trained with fewer features—with a final MAE of 74.29 W/m^2 . This supports our initial hypothesis that the noise model accounts for both the dropped features and other unmeasured and/or unknown disturbances and variables.

Most of the identified features (with exception to the photometers which measure dispersed aerosols) can be obtained with historical measurements, the clear sky model, and a low-cost sky-camera. Notably, two photometer measurements of aerosols are within the top 10 results. Expensive equipment such as a photometer is not likely to be widely used in practice. When the model is retrained with the 8 remaining

Table 5

P-values associated with a paired t-test comparing the validation error of model predictions with and without the noise model. The noise model consistently reduces error across all three steps of model refinement for models with reduced feature counts.

Features	P-value (Noise model reduces error)		
	Step 1	Step 2	Step 3
All	$9.41\text{e}-1$	$3.34\text{e}-1$	$8.24\text{e}-1$
Top 10	$1.08\text{e}-7$	$1.44\text{e}-15$	$1.22\text{e}-1$

inputs, it achieves a MAE of 77.36 W/m^2 , a sacrifice of only 3.07 W/m^2 compared to the 10-feature model. This observation could influence equipment selection at sites interested in irradiance forecasting.

5. Discussion

The final model achieves a MAE of 74.29 W/m^2 , while the POC model has an overall MAE of 134.35 W/m^2 . The distribution of predictions and true values of the final model is shown in Fig. 10 with the persistence model shown below for comparison. These graphs display the normalized density of predicted and true irradiances such that the integrated density in each plot is equal to one. The dashed black line shows a perfect forecast; higher concentrations of the plotted prediction values near this line represent better models. The proposed model reduces the gap between true and predicted irradiance compared to the persistence model. However, the distributions become similar as the prediction horizon increases.

Individual predictions are shown for four sample days in Fig. 11. Notably, POC and the proposed model perform similarly in environments with low irradiance (top left) or relatively constant environments (top right). On days with large ramping events in irradiance (bottom) the proposed model better anticipates changes in weather regime (for example, from sunny to partially cloudy or overcast), though sometimes these predictions are too aggressive.

Fig. 12 shows missing data from the ASI-16 sky camera from outages or known firmware issues. Sky camera images are missing after about 17:00 MST in much of the dataset due to a firmware bug. The missing data may contribute to poor model performance, and may cause significant biases in model performance. Our approach was to only use segments of data without missing data-points. Other works using the NREL SRRL BMS such as Gao and Liu [56], Feng et al. [19] and Al-Lahham et al. [64] also selected only complete sequences. Future work could focus on devising more robust models that account for missing data from the sky camera.




Fig. 10. Mean Absolute Error (W/m^2) distribution the model and POC. As expected, the error increases as time from the forecast (the time horizon) also increases. Density near the correct prediction is notably higher compared to the persistence of cloudiness model. MAE across all predicted intervals is 74.29 W/m^2 .




Fig. 11. True irradiance and the predicted irradiance using the final model and POC for four sample days in the final validation set.

Fig. 13 shows model prediction MAE boxplots for three distinct sky cover conditions at the time of forecast. The values were computed using the validation data. “Clear” indicates CDOC Total Cloud Cover lower than 20%, “Partially Cloudy” indicates CDOC Total Cloud Cover between 20% and 80% and “Overcast” indicates CDOC Total Cloud Cover over 80%. Intuitively, clear conditions are more easily predictable as they follow the persistence of cloudiness assumption. The MAE values for over half of the clear sky condition forecasts were small. However, the model may not anticipate changes in sky conditions and the accompanying changes in clear sky index. This is shown by the large volume of high error outliers in the clear sky category. Not unexpectedly, prediction accuracy drops for partially covered and overcast skies. We explain this by considering the naturally higher likelihood of a change in conditions over the forecast horizon from the time of prediction.

5.1. Comparison to literature

Table 6 presents three selected works that use data collected at the same NREL SRRL BMS site to predict irradiance. We emphasize again that the models proposed in the literature require sequences of images as inputs (and thus have high data transmission requirements) while the model proposed in this work only uses the features extracted from the images, and is thus data-parsimonious (in some cases by several orders of magnitude). Feng et al. [19] use a Sequential CNN model to produce irradiance forecasts from sky camera image sequences. Al-Lahham et al. [64] use clustering on individual image segments to produce irradiance forecasts. Gao and Liu [56] use a transformer model to predict the difference between measured irradiance and the clear sky irradiance.

Table 7 reports RMSE of this work, POC, Feng et al. [19] and Al-Lahham et al. [64]. The model presented in this work has a higher RMSE than the models presented in the literature. However, Fig. 14

Table 6

Selected works using the NREL SRRL BMS sky camera data to predict irradiance. Input sizes assume double types for most data formats and unsigned integer formats for the images.

Reference	Input	Input size (bytes)	Model info
Feng et al. [19]	2 Images	32,768	Sequential CNN
Al-Lahham et al. [64]	6 Images & Aux. Data	14,156,256	Clustering & Random forest
Gao and Liu [56]	8 Images	2,433,024	Visual Transformer, Clear Sky Residual
This Work	10 tabular features	80	CNN-LSTM




Fig. 12. ASI-16 Images missing in the NREL SRRL BMS dataset. A and B were firmware issues noted in the documentation online while C notes a known bug where the camera does not take pictures near 17:00 MST (00:00 GMT).




Fig. 13. MAE of validation forecasts categorized by three distinct cloud cover conditions: Clear indicates CDOC Total Cloud Cover < 20%, Partially Cloudy indicates 20% < CDOC Total Cloud Cover < 80% and Overcast indicates CDOC Total Cloud Cover < 80%. In addition, the star marker for each cloud cover condition indicates the fraction of predictions within each category.

shows a higher FSS for the model proposed in this work compared to the model proposed by Feng et al. [19]. The FSS is a measure of the model's performance relative to the POC model. While the dates of the validation data are not the same, the FSS indicates that the data-parsimonious model proposed in this work may anticipate

changes in weather regime better than more sophisticated models. Interestingly, the FSS of Feng et al. [19] increases over time suggesting that there is some relevant information within the sky camera image sequences used in their model that is not captured in the inputs of the data-parsimonious model.

Integrating the clear sky model and the persistence assumption explicitly into the forecast by changing the representation of future irradiance improves forecasts more than simply providing the clear sky irradiance. These representations simplify model training by isolating the unknown weather effects and not relying on the model to learn already-known long-term dynamics (i.e., de-trending). This work achieves a lower nMAP on the validation period as shown in Table 8 while using a dramatically simpler model.

6. Conclusions

We developed a data-parsimonious machine learning model for near-term forecasting of solar irradiance. The model relies on a CNN-LSTM architecture that includes a noise signal input to account for random and/or unmeasured variables and disturbances that influence irradiance. The model was iteratively refined using three expanding train-validate data splits in order to determine the optimal irradiance representation, time representation, input sequence length, and feature set. The final model achieves a MAE of 74.29 W/m² over a two-hour prediction horizon. While these results have a higher MAE than the models presented in the literature, this model requires orders of magnitude lower data transmission, and has a better FSS indicating that it anticipates changes in weather regime better than the models presented in the literature.

Our findings indicate that predicting the deviation of irradiance from a long-term baseline (e.g., the POC prediction) benefits from a de-trending effect and is thus more accurate than predicting the irradiance itself. Experiments suggested that the optimal forecast only used the most recent measurement of input data rather than a more extensive sequence of measurements. We also found that including a noise model inspired by control theory leads to consistently higher prediction accuracy—even after eliminating all but the most-impactful features. By reducing the input dimensionality, the model is more robust to overfitting on training data and more likely to discern patterns in the data. We expect that these findings will allow for more accurate irradiance forecasting—particularly for locations where data transmission capacity is limited. The findings on the importance of directly integrating the clear sky model and the persistence assumption into the forecast are directly applicable to existing models in the literature and may prove a simple step to improve their performance. Similarly, the noise model may be a useful addition to existing models to account for unmeasured disturbances.

The results presented were limited to 2017–2023 at the NREL SRRL BMS. Future work could expand the dataset to include more years as well as more locations to determine the generalizability of these observations. As models in the literature become open source, future work may also compare the models directly with the same training data rather than only comparing reported results. The irradiance representations and noise model could also be studied across multiple modeling approaches and model structures to understand their individual contribution to irradiance forecasting. Future work could also

Table 7
RMSE (W/m^2) of this work, POC, and reported errors from recent literature.

Forecast horizon (Min.)	POC	This work	Feng et al. [19]	Al-Lahham et al. [64]
10	264.56	102.00	71.30	–
20	254.14	114.68	98.53	–
30	258.72	123.76	109.33	–
40	260.19	130.38	119.35	–
50	264.17	135.91	127.49	–
60	263.40	140.45	135.43	116.7
70	263.86	142.84	–	–
80	269.37	145.95	–	–
90	268.16	150.79	–	–
100	266.79	153.54	–	–
110	265.10	156.20	–	–
120	265.33	158.13	–	127.6




Fig. 14. FSS for a data-parsimonious model and the model presented by Feng et al. [19].

Table 8
nMAP (W/m^2) of POC, this work, and the best results from Gao and Liu [56].

Forecast horizon (Min.)	POC	This work	Gao and Liu [56]
10	21.5	10.2	–
20	22.8	12.1	–
30	24.6	13.6	–
40	26.2	14.6	–
50	27.8	15.6	–
60	29.1	16.4	17.4
70	30.1	16.9	–
80	31.6	17.4	–
90	32.6	18.1	–
100	33.4	18.7	–
110	34.1	19.1	–
120	34.9	19.5	20.9

extend to onsite preprocessing to reduce the need for data transmission by transforming the images into trained embeddings with reduced dimensionality. Finally, the effect of missing data on model performance could be studied further to determine if the model robustness can be increased.

Funding source

Financial support of this work was provided by the Robert A. Welch Foundation (F-1464), the Center for a Solar Powered Future

(SPF2050)—an Industry-University Cooperative Research Center (IU-CRC) funded by the National Science Foundation, United States (EEC-2052814).

CRedit authorship contribution statement

Joshua E. Hammond: Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Ricardo A. Lara Orozco:** Writing – review & editing, Validation, Software, Methodology. **Michael Baldea:** Writing – review & editing, Supervision, Methodology, Formal analysis, Conceptualization. **Brian A. Korgel:** Writing – review & editing, Supervision, Resources, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The original data measurements used for this study were provided by Stoffel and Andreas [53]. Code and dataset used for data aggregation, modeling, and analysis is available at <https://github.com/joshuaeh/TabularSolarForecast>.

Acknowledgments

The authors would like to thank Drs. Herie Soto, Dhruv Arora, Hadi Jamali-Rad, Pierre Carrette and Ojas Shreikar for their feedback in the course of elaborating these results.

Appendix A. Model performance summary

See [Table A.9](#)

Appendix B. All features studied

The features used in this study are organized by source, with units shown in brackets when applicable²:

BMS meteorological station features:

- 315 nm Photometer [nA]
- 400 nm Photometer [μ A]
- 500 nm Photometer [μ A]
- 675 nm Photometer [μ A]
- 870 nm Photometer [μ A]
- 940 nm Photometer [μ A]
- 1020 nm Photometer [μ A]
- Snow Depth [cm]
- Precipitation [mm]
- Precipitation (Accumulated) [mm]
- Station Pressure [mBar]
- Tower Dry Bulb Temperature [deg C]
- Tower Relative Humidity [%]
- Snow Depth Quality [%]
- Station Dry Bulb Temp [deg C]
- Station Relative Humidity [%]
- Vertical Wind Shear [1/s]
- Average Wind Speed at 22 ft [m/s]
- Average Wind Direction at 22 ft [deg from N]
- Peak Wind Speed at 22 ft [m/s]
- Albedo (CM3)
- Albedo (LI-200)
- Albedo Quantum (LI-190)
- Broadband Turbidity
- Sea-Level Pressure (Est) [mBar]
- Tower Dew Point Temp [deg C]
- Tower Wet Bulb Temp [deg C]
- Tower Wind Chill Temp [deg C]
- Airmass
- GHI [W/m^2]
- DNI [W/m^2]
- DHI [W/m^2]

Sky camera features:

- Blue-Red/Blue-Green Total Cloud Cover [%]
- Cloud Detection and Opacity Correction Total Cloud Cover [%]
- Cloud Detection and Opacity Correction Thick Cloud Cover [%]
- Cloud Detection and Opacity Correction Thin Cloud Cover [%]
- Haze Correction Value
- Blue/Red minimum
- Blue/Red median
- Blue/Red maximum
- Apparent Solar Zenith Angle [deg]

- Apparent Solar Azimuth Angle [deg]
- Flag: Sun not visible
- Flag: Sun on clear sky
- Flag: Parts of sun covered
- Flag: Sun behind clouds, bright dot visible
- Flag: Sun outside view
- Flag: No evaluation

Clear sky model features:

- Clear Sky GHI [W/m^2]
- Clear Sky DNI [W/m^2]
- Clear Sky DHI [W/m^2]
- Solar Eclipse Shading
- Zenith Angle [deg]
- Solar Elevation Angle [deg]
- Solar Azimuth Angle [deg]

Engineered features:

- Time from sunrise [Days]
- Time to solar noon [Days]
- Time to sunset [Days]
- Cosine time from sunrise [Days]
- Sine time from sunrise [Days]
- Cosine time to solar noon [Days]
- Sine time to solar noon [Days]
- Cosine time to sunset [Days]
- Sine time to sunset [Days]
- Flag: Day
- Flag: Before solar noon
- Cosine zenith angle
- Cosine normal irradiance
- Wind North-South Speed [m/s]
- Wind East-West Speed [m/s]
- Sun North-South Position
- Sun East-West Position
- Time of Day [Days]
- Time of Year [Years]
- Sine Time of Year [Years]
- Cosine Time of Year [Years]
- Sine Time of Day [Days]
- Cosine Time of Year [Days]
- Clear Sky Index GHI
- Clear Sky Index DNI
- Clear Sky Index DHI
- GHI_{t-1} [W/m^2]
- DNI_{t-1} [W/m^2]
- DHI_{t-1} [W/m^2]
- GHI_{t-2} [W/m^2]
- DNI_{t-2} [W/m^2]
- DHI_{t-2} [W/m^2]
- GHI_{t-3} [W/m^2]
- DNI_{t-3} [W/m^2]
- DHI_{t-3} [W/m^2]
- GHI_{t-4} [W/m^2]
- DNI_{t-4} [W/m^2]
- DHI_{t-4} [W/m^2]
- GHI_{t-5} [W/m^2]
- DNI_{t-5} [W/m^2]
- DHI_{t-5} [W/m^2]
- GHI_{t-6} [W/m^2]
- DNI_{t-6} [W/m^2]
- DHI_{t-6} [W/m^2]

² More information on the NREL SRRL BMS measurements is available at <https://midcdmz.nrel.gov/apps/html.pl?site=BMS;page=instruments>.

Table A.9

Overall model performance summary.

Forecast horizon (Min.)	MAE (W/m ²)	nMAP (%)	RMSE (W/m ²)	FSS (MAE) (%)	FSS (RMSE) (%)
10	45.23	10.16	102.00	54.57	57.01
20	54.69	12.12	114.68	48.07	49.83
30	62.13	13.61	123.76	45.41	47.09
40	67.43	14.63	130.38	44.26	44.99
50	72.45	15.61	135.91	43.72	44.05
60	76.59	16.41	140.45	43.05	42.86
70	79.28	16.92	142.84	42.99	42.46
80	81.84	17.44	145.95	44.00	43.00
90	84.82	18.08	150.79	43.68	41.76
100	87.45	18.68	153.54	43.40	41.54
110	89.13	19.11	156.20	43.45	41.30
120	90.42	19.51	158.13	43.96	41.48
Overall	74.29	16.03	138.89	45.05	44.78

- GHI_{t-7} [W/m²]
- DNI_{t-7} [W/m²]
- DHI_{t-7} [W/m²]
- GHI_{t-8} [W/m²]
- DNI_{t-8} [W/m²]
- DHI_{t-8} [W/m²]
- GHI_{t-9} [W/m²]
- DNI_{t-9} [W/m²]
- DHI_{t-9} [W/m²]
- Clear Sky Deviation GHI_t [W/m²]
- Clear Sky Deviation DNI_t [W/m²]
- Clear Sky Deviation DHI_t [W/m²]
- Mean Clear Sky Deviation $GHI_{t-10:t}$ [W/m²]
- Mean Clear Sky Deviation $DNI_{t-10:t}$ [W/m²]
- Mean Clear Sky Deviation $DHI_{t-10:t}$ [W/m²]
- Median Clear Sky Deviation $GHI_{t-10:t}$ [W/m²]
- Median Clear Sky Deviation $DNI_{t-10:t}$ [W/m²]
- Median Clear Sky Deviation $DHI_{t-10:t}$ [W/m²]
- Clear Sky Index Standard Deviation $GHI_{t-10:t}$ [W/m²]
- Clear Sky Index Standard Deviation $DNI_{t-10:t}$ [W/m²]
- Clear Sky Index Standard Deviation $DHI_{t-10:t}$ [W/m²]
- Mean Clear Sky Deviation $GHI_{t-60:t}$ [W/m²]
- Mean Clear Sky Deviation $DNI_{t-60:t}$ [W/m²]
- Mean Clear Sky Deviation $DHI_{t-60:t}$ [W/m²]
- Median Clear Sky Deviation $GHI_{t-60:t}$ [W/m²]
- Median Clear Sky Deviation $DNI_{t-60:t}$ [W/m²]
- Median Clear Sky Deviation $DHI_{t-60:t}$ [W/m²]
- Clear Sky Index Standard Deviation $GHI_{t-60:t}$ [W/m²]
- Clear Sky Index Standard Deviation $DNI_{t-60:t}$ [W/m²]
- Clear Sky Index Standard Deviation $DHI_{t-60:t}$ [W/m²]

References

- [1] A.K. Magnan, H.O. Pörtner, V.K. Duvat, M. Garschagen, V.A. Guinder, Z. Zommers, O. Hoegh-Guldberg, J.P. Gattuso, Estimating the global risk of anthropogenic climate change, *Nat. Clim. Chang.* 11 (2021) 879–885, <http://dx.doi.org/10.1038/s41558-021-01156-w>.
- [2] A. Guterres, *The Sustainable Development Goals Report 2022*, United Nations publication issued by the Department of Economic and Social Affairs, 2022, p. 64.
- [3] J. Biden, Executive order 12008: Tackling the climate crisis at home and abroad, *Fed. Regist.* (2021) 7619–7633, URL <https://www.federalregister.gov/d/2021-02177>.
- [4] P. Denholm, P. Brown, W. Cole, T. Mai, B. Sergi, M. Brown, P. Jadun, J. Ho, J. Mayernik, M. Colin, R. Sreenath, Examining Supply-Side Options To Achieve 100 % Clean Electricity By 2035, Technical Report, National Renewable Energy Laboratory, Golden, Co, 2022, URL <https://www.nrel.gov/docs/fy22osti/81644.pdf>.
- [5] Energy Information Administration (EIA), *Electric Power Monthly with Data for October 2022*, Technical Report December, United States Energy Information Administration, Washington DC, 2022.
- [6] O. Gandhi, W. Zhang, D.S. Kumar, C.D. Rodríguez-Gallegos, G.M. Yagli, D. Yang, T. Reindl, D. Srinivasan, The value of solar forecasts and the cost of their errors: A review, 2024, <http://dx.doi.org/10.1016/j.rser.2023.113915>.
- [7] X. Chen, Y. Du, E. Lim, H. Wen, K. Yan, J. Kirtley, Power ramp-rates of utility-scale PV systems under passing clouds: Module-level emulation with cloud shadow modeling, *Appl. Energy* 268 (2020) 114980, <http://dx.doi.org/10.1016/j.apenergy.2020.114980>.
- [8] E. Spyrou, V. Krishnan, Q. Xu, B.F. Hobbs, What is the value of alternative methods for estimating ramping needs? in: *IEEE Green Technologies Conference 2020-April*, 2020, pp. 159–164, <http://dx.doi.org/10.1109/GreenTech46478.2020.9289786>.
- [9] K. Lappalainen, J. Kleissl, Analysis of the cloud enhancement phenomenon and its effects on photovoltaic generators based on cloud speed sensor measurements, *J. Renew. Sustain. Energy* 12 (2020) <http://dx.doi.org/10.1063/5.0007550>.
- [10] X. Chen, Y. Du, E. Lim, L. Fang, K. Yan, Towards the applicability of solar nowcasting: A practice on predictive pv power ramp-rate control, *Renew. Energy* 195 (2022) 147–166, <http://dx.doi.org/10.1016/j.renene.2022.05.166>.
- [11] S.C. Rowe, I. Hischer, A.W. Palumbo, B.A. Chubukov, M.A. Wallace, R. Viger, A. Lewandowski, D.E. Clough, A.W. Weimer, Nowcasting, predictive control, and feedback control for temperature regulation in a novel hybrid solar-electric reactor for continuous solar-thermal chemical processing, *Sol. Energy* 174 (2018) 474–488, <http://dx.doi.org/10.1016/j.solener.2018.09.005>.
- [12] V. Bone, J. Pidgeon, M. Kearney, A. Veeraragavan, Intra-hour direct normal irradiance forecasting through adaptive clear-sky modelling and cloud tracking, *Sol. Energy* 159 (2018) 852–867, <http://dx.doi.org/10.1016/j.solener.2017.10.037>.
- [13] B. Cheng, M.D. Collins, Y. Zhu, T. Liu, T.S. Huang, H. Adam, L.C. Chen, Panoptic-deeplab: a simple, strong, and fast baseline for bottom-up panoptic segmentation, *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.* (2020) 12472–12482, <http://dx.doi.org/10.1109/CVPR42600.2020.01249>.
- [14] B. Li, C. Feng, J. Zhang, Multi-timescale simulation of non-spinning reserve in wholesale electricity markets, in: *IEEE Green Technologies Conference 2021-April*, 2021, pp. 520–527, <http://dx.doi.org/10.1109/GreenTech48523.2021.00087>.
- [15] Q. Li, Q. He, Hourly solar irradiance prediction based on enhanced incremental extreme learning machine, in: *International Conference on Optoelectronic Information and Functional Materials, OIFM 2022*, Vol. 1225519, 2022, p. 108, <http://dx.doi.org/10.1117/12.2639376>.
- [16] E.R.C. of Texas, ERCOT Nodal Protocols (3.13 Renewable Production Potential Forecasts), Technical Report, Electric Reliability Council of Texas, 2022, URL <https://www.ercot.com/mktrules/nprotocols/current>.
- [17] Q. Paletta, G. Arbod, J. Lasenby, Omnivision forecasting: combining satellite observations with sky images for improved intra-hour solar energy predictions, 2022, arxiv preprint URL <http://arxiv.org/abs/2206.03207>.
- [18] C. Feng, J. Zhang, Solarnet: A sky image-based deep convolutional neural network for intra-hour solar forecasting, *Sol. Energy* 204 (2020) 71–78, <http://dx.doi.org/10.1016/j.solener.2020.03.083>.
- [19] C. Feng, J. Zhang, W. Zhang, B.M. Hodge, Convolutional neural networks for intra-hour solar forecasting based on sky image sequences, *Appl. Energy* 310 (2022) 118438, <http://dx.doi.org/10.1016/j.apenergy.2021.118438>.
- [20] V. Le Guen, N. Thome, Disentangling physical dynamics from unknown factors for unsupervised video prediction, in: *CVPR*, 2020, pp. 11471–11481, <http://dx.doi.org/10.1109/CVPR42600.2020.01149>.
- [21] A. Dobos, PVWatts Version 5 Manual - Technical Report NREL/TP-6A20-6264, National Renewable Energy Laboratory (NREL), 2014, p. 20, URL <https://www.nrel.gov/docs/fy14osti/62641.pdf>.
- [22] Y. Sun, G. Szűcs, A.R. Brandt, Solar pv output prediction from video streams using convolutional neural networks, *Energy Environ. Sci.* 11 (2018) 1811–1818.
- [23] Y. Sun, V. Venugopal, A.R. Brandt, Short-term solar power forecast with deep learning: Exploring optimal input and output configuration, *Sol. Energy* 188 (2019) 730–741, <http://dx.doi.org/10.1016/j.solener.2019.06.041>.

- [24] L. Clauzel, S. Anquetin, C. Lavaysse, G. Tremoy, D. Raynaud, West african operational daily solar forecast errors and their link with meteorological conditions, *Renew. Energy* 224 (2024) 120101, <http://dx.doi.org/10.1016/j.renene.2024.120101>.
- [25] P.E. Bett, H.E. Thornton, The climatological relationships between wind and solar energy supply in Britain, *Renew. Energy* 87 (2016) 96–110, <http://dx.doi.org/10.1016/j.renene.2015.10.006>.
- [26] P. Niechen, A broadband simplified version of the solis clear sky model, *Sol. Energy* 82 (2008) 758–762, <http://dx.doi.org/10.1016/j.solener.2008.02.009>, URL <https://www.sciencedirect.com/science/article/pii/S0038092X08000406>.
- [27] J.S. Stein, C.W. Hansen, M.J. Reno, *Global Horizontal Irradiance Clear Sky Models: Implementation and Analysis*, Technical Report, Sandia National Laboratories (SNL), Albuquerque, NM, Livermore, CA, 2012.
- [28] M.M. Mikofski, C.W. Hansen, W.F. Holmgren, G.M. Kimball, Use of measured aerosol optical depth and precipitable water to model clear sky irradiance, in: 2017 IEEE 44th Photovoltaic Specialist Conference, PVSC, IEEE, 2017, pp. 110–116.
- [29] R.J. Hyndman, E. Wang, N. Laptev, Large-scale unusual time series detection, in: 2015 IEEE International Conference on Data Mining Workshop, ICDMW, IEEE, 2015, pp. 1616–1619.
- [30] R.J. Hyndman, G. Athanasopoulos, *Forecasting: Principles and Practice*, third ed., OTexts, 2021, URL <http://otexts.com/fpp3>.
- [31] D. Yang, X. Xia, M.J. Mayer, A tutorial review of the solar power curve: Regressions, model chains, and their hybridization and probabilistic extensions, 2024, <http://dx.doi.org/10.1007/s00376-024-3229-4>.
- [32] H. Wen, Y. Du, X. Chen, E.G. Lim, H. Wen, K. Yan, A regional solar forecasting approach using generative adversarial networks with solar irradiance maps, *Renew. Energy* 216 (2023) <http://dx.doi.org/10.1016/j.renene.2023.119043>.
- [33] X. Huang, Q. Li, Y. Tai, Z. Chen, J. Zhang, J. Shi, B. Gao, W. Liu, Hybrid deep neural model for hourly solar irradiance forecasting, *Renew. Energy* 171 (2021) 1041–1060, <http://dx.doi.org/10.1016/j.renene.2021.02.161>.
- [34] G.S. Martins, M. Giesbrecht, Hybrid approaches based on singular spectrum analysis and k- nearest neighbors for clearness index forecasting, *Renew. Energy* 219 (2023) <http://dx.doi.org/10.1016/j.renene.2023.119434>.
- [35] F. Chollet, *Deep Learning with Python*, second ed., Manning Publication, 2021, URL <https://www.manning.com/books/deep-learning-with-python-second-edition>.
- [36] Z. Peng, D. Yu, D. Huang, J. Heiser, S. Yoo, P. Kalb, 3D cloud detection and tracking system for solar forecast using multiple sky imagers, *Sol. Energy* 118 (2015) 496–519, <http://dx.doi.org/10.1016/j.solener.2015.05.037>.
- [37] B. Nouri, P. Kuhn, S. Wilbert, N. Hanrieder, C. Prah, L. Zarzalejo, A. Kazantzidis, P. Blanc, R. Pitz-Paal, Cloud height and tracking accuracy of three all sky imager systems for individual clouds, *Sol. Energy* 177 (2019) 213–228, <http://dx.doi.org/10.1016/j.solener.2018.10.079>.
- [38] S. Song, Z. Yang, H.H. Goh, Q. Huang, G. Li, A novel sky image-based solar irradiance nowcasting model with convolutional block attention mechanism, *Energy Rep.* 8 (2022) 125–132, <http://dx.doi.org/10.1016/j.egy.2022.02.166>.
- [39] F. Lin, Y. Zhang, J. Wang, Recent advances in intra-hour solar forecasting: A review of ground-based sky image methods, *Int. J. Forecast.* 39 (2023) 244–265.
- [40] N. Elizabeth Michael, M. Mishra, S. Hasan, A. Al-Durra, Short-term solar power predicting model based on multi-step CNN stacked LSTM technique, *Energies* 15 (2022) <http://dx.doi.org/10.3390/en15062150>.
- [41] Q. Paletta, A. Hu, G. Arbod, P. Blanc, J. Lasenby, Spin: Simplifying polar invariance for neural networks application to vision-based irradiance forecasting, 2021a, pp. 5182–5191, arxiv preprint <https://arxiv.org/abs/2111.14507v3>.
- [42] Q. Paletta, G. Arbod, J. Lasenby, Benchmarking of deep learning irradiance forecasting models from sky images—an in-depth analysis, *Sol. Energy* 224 (2021b) 855–867.
- [43] G. Liu, H. Qin, Q. Shen, H. Lyv, Y. Qu, J. Fu, Y. Liu, J. Zhou, Probabilistic spatiotemporal solar irradiation forecasting using deep ensembles convolutional shared weight long short-term memory network, *Appl. Energy* 300 (2021) 117379, <http://dx.doi.org/10.1016/j.apenergy.2021.117379>, URL <https://www.sciencedirect.com/science/article/pii/S0306261921007820>.
- [44] S.M.J. Jalali, M. Khodayar, S. Ahmadian, M. Shafie-Khah, A. Khosravi, S.M.S. Islam, S. Nahavandi, J.P. Catalao, A new ensemble reinforcement learning strategy for solar irradiance forecasting using deep optimized convolutional neural network models, in: *SEST 2021-4th International Conference on Smart Energy Systems and Technologies*, 2021, pp. 1–6.
- [45] J. Liu, H. Zang, T. Ding, L. Cheng, Z. Wei, G. Sun, Harvesting spatiotemporal correlation from sky image sequence to improve ultra-short-term solar irradiance forecasting, *Renew. Energy* 209 (2023) 619–631, <http://dx.doi.org/10.1016/j.renene.2023.03.122>.
- [46] E. Ogliari, M. Sakwa, P. Cusa, Enhanced convolutional neural network for solar radiation nowcasting: All-sky camera infrared images embedded with exogenous parameters, *Renew. Energy* 221 (2024) <http://dx.doi.org/10.1016/j.renene.2023.119735>.
- [47] M.S. Ghoniya, B. Urquhart, C.W. Chow, J.E. Shields, A. Cazorla, J. Kleissl, A method for cloud detection and opacity classification based on ground based sky imagery, *Atmos. Meas. Tech.* 5 (2012) 2881–2892, <http://dx.doi.org/10.5194/amt-5-2881-2012>.
- [48] A. Kazantzidis, P. Tzoumanikas, A.F. Bais, S. Fotopoulos, G. Economou, Cloud detection and classification with the use of whole-sky ground-based images, *Atmos. Res.* 113 (2012) 80–88, <http://dx.doi.org/10.1016/j.atmosres.2012.05.005>.
- [49] A. Heinle, A. Macke, A. Srivastav, Automatic cloud classification of whole sky images, *Atmos. Meas. Tech.* 3 (2010) 557–567, <http://dx.doi.org/10.5194/amt-3-557-2010>.
- [50] S.A. Haider, M. Sajid, H. Sajid, E. Uddin, Y. Ayaz, Deep learning and statistical methods for short- and long-term solar irradiance forecasting for Islamabad, *Renew. Energy* 198 (2022) 51–60, <http://dx.doi.org/10.1016/j.renene.2022.07.136>.
- [51] D. Makwana, S. Nag, O. Susladkar, G. Deshmukh, S.C. Teja R, S. Mittal, C.K. Mohan, Aclnet: an attention and clustering-based cloud segmentation network, *Remote Sens. Lett.* 13 (2022) 865–875, <http://dx.doi.org/10.1080/2150704X.2022.2097031>.
- [52] Y. Fabel, B. Nouri, S. Wilbert, N. Blum, D. Schnaus, R. Triebel, L.F. Zarzalejo, E. Ugedo, J. Kowalski, R. Pitz-Paal, Combining deep learning and physical models: a benchmark study on all-sky imager-based solar nowcasting systems, *Solar RRL* (2023) <http://dx.doi.org/10.1002/solr.202300808>.
- [53] T. Stoffel, A. Andreas, Nrel Solar Radiation Research Laboratory (Srrl): Baseline Measurement System (Bms); Golden, Colorado (Data), Technical Report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 1981, <http://dx.doi.org/10.5439/1052221>.
- [54] EKO, Asi-16 find clouds software, 2018, URL <https://www.eko-instruments.com/media/z2aalsq/asi-16-software-manual-find-clouds.pdf>.
- [55] W.F. Holmgren, C.W. Hansen, M.A. Mikofski, Pvlb python: a python package for modeling solar energy systems, *J. Open Sour. Softw.* 3 (2018) 884, <http://dx.doi.org/10.21105/joss.00884>.
- [56] H. Gao, M. Liu, Short-term solar irradiance prediction from sky images with a clear sky model, in: *Proceedings - 2022 IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2022*, 2022, pp. 3074–3082, <http://dx.doi.org/10.1109/WACV51458.2022.00313>.
- [57] J.B. Rawlings, D.Q. Mayne, M.M. Diehl, S. Barbara, *Model Predictive Control: Theory, Computation, and Design*, second ed., 2022, URL <http://www.nobhillpublishing.com>.
- [58] P. Bauer, Multiple testing in clinical trials, *Stat. Med.* 10 (1991) 871–890.
- [59] A. Altmann, L. Toloş, O. Sander, T. Lengauer, Permutation importance: a corrected feature importance measure, *Bioinformatics* 26 (2010) 1340–1347.
- [60] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2014, arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [61] I. Loshchilov, F. Hutter, Sgdr: Stochastic gradient descent with warm restarts, 2017, <http://arxiv.org/abs/1608.03983>.
- [62] A.M. Stolwijk, H. Straatman, G.A. Zielhuis, Studying seasonality by using sine and cosine functions in regression analysis, *J. Epidemiol. Community Health* 53 (1999) 235–238, <http://dx.doi.org/10.1136/jech.53.4.235>.
- [63] R. Bellman, *Dynamic programming*, *science* 153 (1966) 34–37.
- [64] A. Al-Lahham, O. Theeb, K. Elalem, T.A. Alshaw, S.A. Alshebeili, Sky imager-based forecast of solar irradiance using machine learning, *Electron. (Switz.)* 9 (2020) 1–14, <http://dx.doi.org/10.3390/electronics9101700>.