

A deep learning model for intra-day forecasting of solar irradiance using satellite-based estimations in the vicinity of a PV power plant

Emilio Pérez^{a,*}, Javier Pérez^b, Jorge Segarra-Tamarit^a, Hector Beltran^a

^aDepartment of Industrial Systems Engineering and Design, Universitat Jaume I, Castelló de la Plana, Spain.

^bInstituto Tecnológico de Informática (ITI), Universidad Politécnica de Valencia, Valencia, Spain

Abstract

This work proposes an intra-day forecasting model, which does not require to be trained or fed with real-time data measurements, for global horizontal irradiance (GHI) at a given location. The proposed model uses a series of time-dependant irradiance estimates near the target location as the main input. These estimates are derived from satellite images and are combined with other secondary inputs in an advanced neural network, which features convolutional and dense layers and is trained using a deep learning approach. For the various input combinations, the performance of the model is validated with a quantitative analysis on the forecast accuracy using different error metrics. Accuracies are compared with a commercial solution for irradiance forecasting made by the European Centre for Medium-Range Weather Forecasts (ECMWF) and publications with similar approaches and forecasting horizons, showing state-of-the-art performance even without irradiance measurements.

Keywords: Irradiance forecasting, deep learning, neural networks, satellite data.

1. Introduction

Information on solar irradiance has been a major interest for meteorologists for centuries. In many recently published studies (Yang et al., 2018), irradiance forecasting achieves precise and accurate results owing to a variety of developing technologies. In fact, photovoltaic (PV) is a crucial, fast growing technology that has established itself as a pivotal player internationally by continuously increasing its share in the global power generation industry. The continuous growth of this industry in the last 10-12 years, with more than 100 GW of newly grid-connected capacity in 2018, has led to a total current capacity above 500 GW worldwide (Jaeger-Waldau, 2018). Therefore, such a significant number of current and expected PV deployments implies elevated levels of PV penetration in multiple electric power systems. This casuistic menaces the grid stability (Katiraei & Aguero, 2011) due to the inherently intermittent nature of PV production, caused by passing clouds. This prevents PV power plants from achieving accurate predictions for future productions, making it difficult for grid operators and PV owners to manage the system and deal their production in the electricity market. Such limitations have been traditionally dealt by integrating the production of multiple PV plants that are geographically dispersed (Marcos et al., 2012), and more recently, by introducing different types of energy storage systems (mainly batteries) next to the PV plant (Hanna et al., 2014). However, batteries still represent a significant cost factor for PV plants and suffer from accelerated ageing if the operational conditions are not well controlled (Jossen et al., 2004).

In this context, the possibility to forecast the irradiance and, hence, the PV production, becomes a fundamental factor. Irradiance forecasting increases PV reliability (dispatchability) (Tuohy et al., 2015). In addition, when batteries are introduced in the PV plant, it allows minimizing the battery size required to operate such a hybrid power plant reliably in the electricity market (Brenna et al., 2016). Finally, the irradiance forecast allows optimizing the plant operation, profiting information regarding future production, and thus minimizing the battery ageing.

Several types of solar forecasting techniques exist. Based on the approach used, they can be broadly categorized into two main groups (Ren et al., 2015): physical approaches, which use knowledge based on atmospheric science, and data-driven approaches, which use historical data as the main input for prediction.

In order to choose a correct approach it is essential to consider the target forecasting horizon and the time-step, or granularity, which depend on the projected use of the forecast (Voyant et al., 2017). This study focuses on intra-day forecasts (with a horizon up to 6 h ahead). For such a forecasting horizon, the prevailing techniques use satellite data, sometimes combined with real-time measurements (Blaga et al., 2019) and ensemble methods, which use multiple predictors to obtain an aggregated improved decision, with regard to that provided by the base predictors (Ren et al., 2015). For shorter forecasting horizons, time series and sky images are used (Zhao et al., 2019), whereas for longer forecasting horizons, numerical weather predictions are preferred (ECMWF, 2016).

Different reviews in the literature refer to approaches based on machine learning methods (Voyant et al., 2017) such as artificial neural networks (ANNs) (Qazi et al., 2015). Both reviews conclude these methods offer accurate prediction results,

*pereze@uji.es

although comparisons are always difficult to perform because of the varying locations and conditions for the different forecasts. While performing the literature review on recent solar forecasting studies, such comparability issues, along with some reproducibility issues, are very common. To avoid this, we will try to follow the ROPES guidelines for the solar forecasting proposed in (Yang, 2019).

Many works that implement ANNs, such as (Mazorra Aguiar et al., 2015), use ground measurements and satellite data for forecasting, achieving relative root-mean-square error (rRMSE) values from 15.3 % to 36.2 %, depending on the horizon and variability of the location. In (Blaga et al., 2019), the performance of different forecasting methods is classified by its forecasting horizon and the climatology of the location. For temperate climate and intra-day forecasts, the rRMSE of the methods varies from 20 % to 40 %, with machine learning and cloud motion methods having the best performance. Machine learning methods can also be found in (Lauret et al., 2015) with errors between 20 % and 25 % for 1 h ahead forecasts. Moreover, cloud motion approaches achieve errors for 1 h ahead forecasts of approximately 28 % in (Dong et al., 2014) and 10 % in (Alonso-Montesinos & Batlles, 2015) (dividing by the difference between the maximum and the minimum measured irradiance instead of the mean). Authors in (Gutierrez-Corea et al., 2016) also propose an ANN approach that uses nine input sequences of 10 values from 10 different locations for forecasting. The results using this method show errors between 22.6 % and 32.1 % for 1-h and 6-h forecasting horizons.

Note that almost all of these works share two common features: they use satellite images as a direct input to their forecast models as well as real-time irradiance measurements. The former can pose a problem because it burdens the forecast model not only by determining how the meteorological conditions will evolve in time but also with how to compute the irradiance from satellite images. The latter implies the use of information that in many occasions is not reliable, or not even available. Small installations with budget constraints, remote locations, and facilities with difficulties for appropriate maintenance of data acquisition systems are some case examples.

Some other works in the literature have attempted to solve these limitations. For instance, (Larson & Coimbra, 2018) uses satellite images and a support vector regression (SVR) to forecast irradiance without the use of real-time measurements. However, actual measurements are still needed to fit the SVR model. Authors in (Lago et al., 2018) go one step beyond to address both these limitations by developing a deep neural network (DNN) whose inputs are: the numerical weather prediction (NWP) for the location, the clear sky irradiance, and a satellite derived GHI for the target location, which is based on the same empirically adjusted algorithm later described in this work. Their model can be trained with actual data from a few locations and deployed in places where no telemetry is available. Although these two works attempt to get around the use of measured irradiance, their models still need to be trained with measured data.

In order to overcome these drawbacks, the DNN developed in this work forecasts intra-day irradiance using a deep learn-

ing structure, which does not require real-time data measurements. Furthermore, the proposed model uses an estimation of the past GHI, which is not only in the target location but also in its vicinity, as the main input. These data matrices, which are obtained from a physics-based empirically adjusted online algorithm based on satellite data, are fed into several consecutive convolutional layers to detect features and perform a type of deep-learning-based cloud motion detection. This information is then combined with other secondary inputs and fed into the DNN dense layers to obtain the final irradiance forecast. Therefore, even in the absence of measured irradiance data, the proposed DNN can obtain the irradiance forecasts with state-of-the-art performance.

In the next section, the input datasets used for the proposed deep learning approach as well as the structure of the DNN itself are introduced. Section 4 summarizes the forecasting results for the different situations and available datasets analysed. Finally, the conclusions and findings are drawn.

2. Model inputs

Solar irradiance in a given location at a given time depends on a number of factors that can be split into two groups, those that are independent from the atmosphere state and those that are not. The first group, which includes the hour angle in the local solar time, current declination of the Sun, and local latitude, are usually taken into account with a single parameter, the solar zenith angle (θ_s), which precisely determines the solar irradiance in a horizontal plane on the top-of-atmosphere (TOA), I_{TOA} .

Regarding the parameters that depend on the atmospheric state, they can be further split into two different types, those that determine the irradiance in a clear-sky case (I_{cs}) and those that quantify the depletion of I_{cs} due to clouds (Wald et al., 2015).

I_{cs} depends, mainly, on elements in the atmosphere such as aerosols, water vapour, and ozone content (Wald et al., 2015) and can be calculated using several models, such as the Bird model (Bird et al., 1986).

In actual conditions, I_{cs} is a challenging value to estimate and even more difficult to forecast. Furthermore, the reduction of I_{cs} due to clouds, which changes more rapidly, makes the estimation and forecasting of GHI on the surface a harder task.

Therefore, instead of trying to separately forecast all these factors, the main idea for our irradiance forecasting model includes inferring changes associated to the movement of the clouds by using a series of consecutive datasets of the available estimated irradiance centered in a target location. Fig. 1 shows the hourly datasets of the estimated irradiances in a given region. Analysing the change between the first two images, it can be deduced by the irradiance values that there is a cloud around the upper-left corner of the image advancing towards the centre and, therefore, the third image seems logical. This is the type of information we expect our DNN to infer.

2.1. Past solar irradiance estimation

From the previous discussion, it becomes clear that the main input for the irradiance forecast is the estimated irradiance in




Figure 1: Hourly evolution of irradiance centred in the target location.

the surroundings of the target location. There are multiple ways to obtain such an estimation (Zhang et al., 2017); some of them are even based on ANNs (Quesada-Ruiz et al., 2015). However, because of the convenience in obtaining results for an arbitrarily large region with a good time granularity in real time, this work uses datasets obtained by the Surface Insolation under Clear and Cloudy Skies (SICCS) algorithm (Greuell et al., 2013). This is a physics-based empirically adjusted algorithm developed for estimating the surface solar irradiance from satellite data. The most important inputs are a cloud mask product and the cloud properties dataset derived from Meteosat/Spinning Enhanced Visible and Infrared Imager (SEVIRI) observations. These observations are given with a temporal resolution of 15 min and a nadir spatial resolution of $3 \times 3 \text{ km}^2$. Using these irradiance datasets as inputs helps to avoid dealing with multiple satellite channels to estimate the properties of different cloud layers and, therefore, the convolutional layers of the DNN can focus on the cloud motion detection.

These estimates are available for all the spatial domain of Meteosat Second Generation (longitude $[-50, 50]$, latitude $[-80, 80]$), which covers Europe and Africa as well as small parts of Asia and America. However, the forecast model proposed in this work can also be used outside this domain using previously implemented SICCS algorithm with images from other satellites.

Archives of GHI estimates are available at <http://msgcpp.knmi.nl>, which contain the last three years of data. However, note that these GHI matrices are not available instantly as there is a 20-min delay until the current information is downloadable.

2.2. TOA/Clear-sky solar irradiance forecast

Although the main input for the forecasting DNN is the described series of irradiance estimates, some other information sources are also used.

The I_{TOA} forecast, as previously discussed, can be precisely computed when the exact location and time are known, as in this case. This information can help in prediction by providing the shape of the maximum irradiance potentially available.

However, if an estimation of the future I_{cs} could be fed into the DNN instead of the I_{TOA} , it would only need to forecast the cloud-related phenomena. In general, I_{cs} will not be exactly known in advance, but there are available forecasts on many of the factors such as aerosols and ozone that affect it. In particular, the Copernicus European Union Programme provides, in its Atmosphere Monitoring Service (CAMS), forecasts of some of

these parameters such as the different types of aerosol concentrations, ozone, and other greenhouse gases (ECMWF, 2016). Therefore, some of them could be used as inputs to help the DNN with the forecast. However, all those forecasts are still not used as inputs in this work. Instead, I_{TOA} will be substituted with an estimation of I_{cs} obtained from the Copernicus service. Although this approach would not be possible in real-time applications, it provides preliminary results that allow us to quantify the maximum improvement achievable with that information and, therefore, decide whether they are worth including in the network structure.

2.3. Measured irradiance feedback

Although one of the main advantages of the proposed model is the ability to provide a forecast without the need of irradiance measurements, it is convenient to quantify to what extent the forecasting errors can be reduced if the actual on-site measurements of irradiance, obtained in the target location using a pyranometer, were available.

First, it is interesting to analyse the effect of having a good amount of actual irradiance data from the system to train the DNN. These data should cover all the possible system conditions. Therefore, they should span for more than a year to split them to be used as training, validation, and test data.

However, the existence of the pyranometer data for the training stage does not necessarily imply that those measurements will be continuously available during the normal operation of the system. Depending on the nature of the PV system making use of the prediction, in some cases there will be no need for such an equipment (as in domestic installations). Subsequently, training and validation data can be acquired with a portable pyranometer during a limited time span. In other cases, the forecast service may not be implemented at the system location, and the irradiance feedback would require bidirectional communication, making the entire system more expensive. Furthermore, pyranometers require a calibration and maintenance service to provide reliable data that is not always carried out.

Therefore, three studies on the forecast performance are introduced in the results section. These include when the irradiance feedback is not available, when it is only available while training, and when it is also available in real time.

3. Model selection and evaluation

As will be discussed later, there are a number of choices that have to be made when defining the architecture of a DNN (regarding the number and type of layers, activation functions, etc...). This is the so-called NN model selection. Typically, this selection is carried out by calculating the performance achieved by some error metric, when considering the different model alternatives. To do so, the two and a half years of available data are split into two different datasets. The first one covers two years of data (2015 and 2016) which are randomly shuffled and used to select the best model architecture. The second dataset, corresponding to the first half of 2017, is used in order to obtain the performance results presented in this work.

The model selection dataset is also split into three different subsets:

- The *training data*: correspond to irradiance values (either real measurements or estimations) registered at the target location during, at least, a whole year. At each iteration of the training, data chunks are shuffled and reintroduced in batches of 80 samples, so that the DNN does not overfit to the original data sequence. Furthermore, to avoid an undesired bias of data and to speed up the training, nights are “shortened” to a few hours. This data split contains 50% of the model selection data.
- The *validation data*: correspond to irradiance values registered during a different (and shorter) time period from the training one. These data are evaluated periodically during the training as a sustained increase of the error is an indication of overfitting of the DNN. Therefore, an early stopping strategy is implemented when this occurs, which allows obtaining a general solution to the problem. As for the *training data*, nights are shortened for this dataset. This data split contains 25% of the model selection data.
- The *test data*: correspond to irradiance values registered during a different time period from that of training and validation data. These data are used to evaluate and compare the performance of the different DNN topologies once their training is completed. Note that all the data with solar zenith angles exceeding 80° ($\theta_s > 80^\circ$) are removed from this dataset. This is because time periods just after the sunrise and before the sunset present low precision pyranometer readings due to the inherent low irradiance levels. This encompasses the nights because forecast is trivial during that part of the day for all the topologies. The performance of the different model architectures are compared using this dataset, which contains the remaining 25% of the model selection data.

Furthermore, it is convenient to avoid reporting as results of the forecasting model those obtained for the test data with the best possible structure, as these results have, to some extent, embedded information from the data during the model selection. In this work, data corresponding to the first half of 2017 are used to evaluate the models and calculate the final results.

GHI measurements from a pyranometer were collected for two and a half years at a PV plant located in the southern part of France, with latitude 43.4° and longitude 6.0° . Furthermore, 151×151 GHI matrices centred in this location were downloaded and stored as daily files with $151 \times 151 \times 96$ values. After downloading the data, 96 days presented missing values and were deleted from the dataset.

At the studied location, summers are hot and relatively dry whereas winters are mild and rainy. The annual global irradiance for the site is approximately 1600 kWh/m^2 , whereas the annual rainfall is approximately 700 mm.

As shown in (Mazorra Aguiar et al., 2015), Fig. 2 analyses the climatology at the target location by representing the clear-sky index, $K_{cs} = I/I_{cs}$ and the standard deviation of its variation

for each day. On the one hand, zones A, B, and C represent how cloudy a day is, with zone A being the cloudiest and zone C being a clear sky. On the other hand, zones I, II, and III show the stability of the cloudiness throughout each day. Therefore, days in zone III are the hardest to predict, and days in the zone CI are the easiest, as they are clear sky days with low variability.




Figure 2: Clear-sky index (K_{CS}) distribution for the target site.

All the GHI measurements used in this work are available at (Segarra-Tamarit et al., 2020) to facilitate the reproducibility of this research. Furthermore, as the archive of GHI estimates at <http://msgcpp.knmi.nl> only stores the last three years of data, the matrices used in this work are also provided within this repository.

3.1. Metrics used for performance evaluation

The performance of the proposed DNN models is tested against real-time data measured using a pyranometer for two reasons. First, during the model selection procedure, in order to quantitatively determine the best possible structure. Second, during the model evaluation, in order to report the final results. Furthermore, in the model evaluation stage, this work explores and quantifies how much a forecast improves when additional datasets are added.

To carry out this evaluation, some of the most accepted standard evaluation metrics (Voyant et al., 2017) are used. These include:

- **Mean absolute error and relative mean absolute error:**

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{I}_i - I_i| \quad (1)$$

$$\text{rMAE} = \frac{\text{MAE}}{\bar{I}}, \quad (2)$$

where I_i and \hat{I}_i are the measured and estimated irradiance values, respectively, N is the number of measurements, and \bar{I} is the mean value of the measurements for the whole period.

- **Root-mean-square error and relative root-mean-square error:**

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{I}_i - I_i)^2} \quad (3)$$

$$\text{rRMSE} = \frac{\text{RMSE}}{\bar{I}} \quad (4)$$

- **Forecast Skill:**

$$\text{Skill} = \left(1 - \frac{\text{RMSE}_{\text{model}}}{\text{RMSE}_{\text{pers}}}\right) \cdot 100, \quad (5)$$

where $\text{RMSE}_{\text{model}}$ and $\text{RMSE}_{\text{pers}}$ are the RMSE calculated according to (4) for the model considered and the smart persistent model, respectively. This smart persistent model used here as a reference, is a naive forecast technique usually seen in the literature that consists of determining the current clearness index, assuming it will remain constant in the future:

$$\hat{I}_{\text{pers},i+t} = K_C(i)I_{\text{TOA},i+t}, \quad (6)$$

with

$$K_C(i) = \frac{I_i}{I_{\text{TOA},i}}. \quad (7)$$

Note from previous equations that the MAE weights linearly forecast errors whereas the RMSE is more sensitive to big errors. This implies that the latter is more suitable for applications where small errors can be tolerable, but large ones cause increased losses. Hence, the suitability of each evaluation metric is heavily dependent on the application.

Relative values for both magnitudes are also calculated when dividing by the mean value, because they are usually more intuitive to understand and allow some comparison with other works. However, it should be considered that different datasets are in general difficult to compare, as the variability of the locations may differ considerably.

Forecast skill is introduced as a simple way of analysing how a given model outperforms the smart persistent model, or any other model used as a reference. Moreover, this metric allows a fair comparison among point forecasts made at different locations and horizons (Yang, 2019).

Finally, when different forecasts are compared, an important parameter that needs to be considered is the data granularity. The lower the time step, the higher the evaluation errors as the forecast model should be able to reproduce fast phenomena. In this work, a base time step of 15 min is used, matching the frequency of the satellite data. However, as many results in the literature are given for a 1-h granularity, a new forecast with this time step is calculated from the 15-min as

$$\hat{I}_{1h,j} = \frac{1}{4} (\hat{I}_{4j} + \hat{I}_{4j-1} + \hat{I}_{4j-2} + \hat{I}_{4j-3}). \quad (8)$$

Then, the irradiance of every hour is calculated as the mean of the last four values of irradiance for the 15-min time step forecast. As we will see, errors with this granularity are lower

because the fastest phenomena, which are more difficult to predict, are filtered by the re-sampling process. However, results could probably be improved by DNNs properly trained for this granularity.

As similar conclusions can be extracted by comparing the different error metrics, only rRMSE and Skill (being the most used metrics in the literature) are listed in Tables 1 and 2 and will be discussed in the results section. However, for the sake of completeness and easier comparison with other works, the rest of error metrics (MAE, rMAE and RMSE) are also included in Tables 4 and 5 in Appendix A.

3.2. Network selection and best structure

For the selection of the best DNN model, various topologies were initially considered and their results compared by means of the test data. A grid search of potential DNNs covered different combinations of hyperparameters such as number of neurons, layers, and activation functions. The training and evaluation of the different topologies were executed using Python 3.6 and Tensorflow 1.14.

Therefore, for each possible topology, the training procedure started with a random initialization of the DNN parameters. Subsequently, the *training data* batches were fed to the model to perform an end-to-end training, i.e. fitting the parameters for every layer simultaneously. An Adam optimizer was used for the training to minimize the MSE, which mainly penalizes big errors, and L2 regularisation. The training stopped once the performance of the DNN worsened several times on the *validation data*.

Among the multiple alternatives analysed, Fig. 3 shows the selected structure. It consists of five convolutional layers followed by three dense layers. The convolutional layers are introduced for extracting features from past estimated-irradiance datasets, which are treated as a series of images that report a value of irradiation for each time instant and for every pixel. The dense layers are used to perform a more general analysis of the inputs and infer the final irradiation forecast.

The main input to the DNN, which is the irradiance estimates of the previous prediction time, is introduced to the initial convolutional layer as 10 image channels, corresponding to 2h30m. Although 151×151 matrices were initially considered, empirical tests with sizes beyond 35×35 did not show significant improvements, because of the inherent behaviour of the meteorological phenomena. Therefore, this input consists of 10 35×35 matrices of irradiance estimates centred at the target location.

The convolutional layers can be compared with the cloud motion methods because they perform mathematical operations exclusively affecting pixels surrounding the one being processed. In this way, it is possible for the convolutional layer to detect shapes or objects invariantly to their position. This enables the DNN to extract features from the pattern of the clouds in the different images and predict their movement or evolution. Out of the five convolutional layers, each with a 5×5 kernel size, the first four present a pooling layer whose strides and size are 2 in all dimensions. Besides, the convolutional layers present a




Figure 3: Proposed deep-learning structure with convolutional layers for cloud motion detection and feed-forward layers to perform 15-min forecasts.

rectified linear unit (RELU) activation. The first layer extracts 30 features, the second layer extracts 40 features, and the other layers extract 50 features each. As the data travels from one layer to the next, its dimension is reduced by pooling layers and more abstract information is obtained.

Subsequently, the output of the convolutional layers, together with the rest of the inputs provided to the DNN, i.e. the TOA or clear-sky irradiation curve and the measured irradiance data, are fed into the three consecutive dense layers. This dense structure, with three fully connected layers of 300, 100, and 1 neurons and also presenting a RELU activation, is replicated in parallel to provide irradiation forecasts from 15 min to 6 h ahead, with a 15-min temporal resolution.

Finally, note that for the operational use of this model some time-related parameters must be considered during the implementation (Yang et al. (2019)). The main issue to be considered is the delay in the last available GHI matrix, which forces a minimum lead-time of 30 min. This implies that the first two forecasts are useless as they are forecasting values already experienced. Moreover, note that all the forecasts for the next 6 h, with a 15-min time resolution, can be updated every 15 min. With these characteristics, the DNN forecasting model can be productively applied at PV power plants integrating energy storage systems and participating in the intra-day electricity markets.

4. Results and discussion

Once the best model structure has been determined using the grid search process, new DNNs with this structure were trained using the same data as for the model selection and tested using the model evaluation data. Furthermore, although the main focus is the GHI forecast without the use of actual data measurements, the other possible cases are also considered here. These involve scenarios in which more information is available in order to quantify the impact on performance of introducing different datasets. All the results are shown in Table 1.

4.1. Top-of-atmosphere Copernicus and actual measured data

The inputs for the DNN of the first considered case include a series of 10 estimated radiation datasets of size 35×35 (approximately $100 \times 100 \text{ km}^2$), TOA irradiance from Copernicus (I_{TOA}), and the last measurement of the irradiance in the location. The DNN is trained with actual data as previously stated. The results for this model are labelled in all tables and figures as *TOA & all real*.

Fig. 4 shows some of the forecasts for the 1-h granularity version of this model for a clear and cloudy day. Actual measurements are represented using a solid line, whereas each dashed line represents the forecast obtained at the time instant indicated by the square markers. It can be seen that for sunny days the model achieves a good forecast. For the cloudy day, it is more challenging in the first few hours of the day (while the satellite still does not provide irradiance estimations), but the model improves afterwards.

Fig. 5 graphically compares how the rRMSE evolves as the prediction horizon increases for this model versus the *persistent* model. Note that both models perform similarly for the near future, but predictions degrade much earlier for the persistent model because of K_c changes.

4.2. Clear-sky model

The objective of the second analysis is to determine if the forecasts can be significantly improved by adding an additional input providing information about the characteristics of the atmosphere, regardless of the cloud cover. With this aim, a new DNN is trained substituting the Copernicus TOA input with the Copernicus *clear-sky* estimation. As previously discussed, it is not realistic to assume this dataset will be available but forecasts for some of the variables that could generate it do exist and could be used as inputs in future developments. This model is labelled as *Clear-sky & all real*.

Fig. 5 shows the rRMSE of the *clear-sky* model used as a forecast on its own, as well as the *Clear-sky & all real* along with the previously discussed *persistent* and *TOA & all real*

Table 1: Forecasting errors for the considered models with a 15-min granularity

		15min	30min	...	1h	...	2h	...	3h	...	4h	...	5h	...	5h45min	6h
rRMSE (%)	Persistent	18.0	22.5	-	27.1	-	36.3	-	44.0	-	49.6	-	54.3	-	56.7	57.2
	TOA & all real	15.6	18.2	-	21.2	-	26.0	-	29.0	-	32.0	-	35.0	-	36.6	37.7
	Clear-sky & all real	15.5	18.2	-	21.3	-	25.1	-	28.7	-	31.6	-	34.4	-	35.9	35.9
	TOA & real labels	17.7	19.4	-	22.4	-	26.4	-	29.9	-	33.1	-	35.6	-	37.1	37.5
	TOA & all satellite	21.8	22.3	-	24.5	-	28.4	-	31.8	-	34.4	-	37.0	-	38.3	38.8
Skill (%)	TOA & all real	13.9	19.5	-	21.8	-	28.5	-	34.3	-	35.8	-	35.8	-	35.8	34.4
	Clear-sky & all real	14.3	19.5	-	21.6	-	31.0	-	35.1	-	36.6	-	36.8	-	37.0	37.5
	TOA & real labels	2.1	13.9	-	17.5	-	27.5	-	32.3	-	33.6	-	34.8	-	34.9	34.7
	TOA & all satellite	-20.9	1.3	-	9.8	-	22.0	-	27.9	-	30.9	-	32.2	-	32.9	32.5




Figure 4: TOA & all real predictions with 1-h granularity for a cloudy and a sunny day as compared with actual irradiance measurements.




Figure 5: Relative RMSE with increasing horizons, 15-min granularity.

models. Note how the rRMSE for the clear-sky estimation model does not change as the horizon moves away because both real data and the prediction remain the same. That is due to the fact that this model does not use irradiance datasets as inputs and, therefore, it does not benefit from being closer in time to the instant it is trying to predict. Thereafter, this model offers poor forecasting results by itself.

Regarding the last added model in Fig. 5, *Clear-sky & all real*, notice how approximately during the first four hours the clear-sky information does not improve the forecast. Nonetheless, from that horizon onwards, its rRMSE tends to be lower than that of the *TOA & all real* model. This is probably because the evolution of past estimated irradiance datasets have the clear sky index information embedded in it, and it changes slowly. Then, as the prediction horizon moves away, this index changes and the model benefits from having that information available.

4.3. Use of actual measurements

Next, we will try to quantify the effect that the use of measured irradiance has on the forecasting error metrics. Our aim is to compare the *TOA & all real* model with two new models:

- *TOA & real labels*: In this case, the actual irradiance is not continuously measured at the installation and therefore, it cannot be used as an input for the DNN. The structure of the DNN is exactly the same as used before, but instead of the actual measurement, the past estimated irradiance value is used. However, historical actual data from the site are considered to be available for the training of the DNN.
- *TOA & all satellite*: This second case assumes there is no availability of actual data, neither in real-time nor in historical series. As in the previous case, actual data are replaced by the estimated irradiance data. However, note that the model is validated with actual data.

Both analyses are useful to perform because if the deterioration of the forecast is acceptable, the cost of implantation of the forecasting model would be reduced. In the first case, a portable pyranometer could be left on the target location giving enough time to acquire the training data and subsequently be used elsewhere, without the need of leaving a permanent one. In the second case, no pyranometers would be required.

It can be concluded from Fig. 6 that, as expected, the rRMSE is lower when actual data is available for training, and even more if it can be measured during operation. However, some other conclusions can also be made. First, the real-time irradiance feedback is slightly more important for the short-term predictions than for longer term prognosis. Second, the forecast achieved when no actual data are used, although worse, still seems useful in many situations.

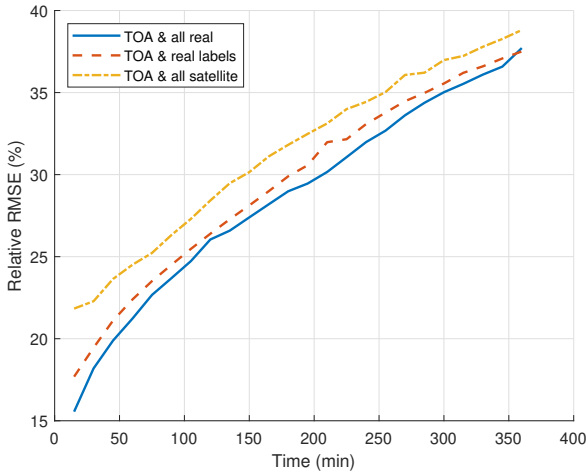


Figure 6: Relative RMSE with increasing horizons, 15-min granularity.

Previous conclusions are confirmed by the forecast skill results shown in Fig. 7. For the very-short term, the skill of *TOA & all satellite* and *TOA & real labels* is very low (even negative for *TOA & all satellite*), making it clear that the knowledge of recent measurements is more important than the satellite information. However, all models outperform the *persistent* model for horizons beyond 30 min. In fact, the higher the horizon, the lower the impact in the forecast skill of actual measurements and the higher the impact of using a clear-sky model. Finally, it can be appreciated how, for horizons ranging from 5h30 onwards, the forecast skill starts to decrease, showing the well-known limitation of satellite-based forecasts for longer horizons.

4.4. One-hour granularity

Finally, results with the 1-h granularity model are listed in Table 2 and shown in Fig. 8. As expected, errors for every horizon and model are always lower than their 15-min counterpart, because of the filtering effect that averaging the 15-min predictions and measurements introduces.




Figure 7: Forecast skill, 15-min granularity.

It is interesting to compare these results with the rRMSE obtained in (Lago et al., 2018). This is the contribution from the literature with the most similar approach to the one proposed here because they also avoid the use of measured irradiance for their forecasts. In that work, rRMSE is 25.07 % for the 1-hour horizon, 38.71 % for the 6-hour horizon and 31.31% for the mean value of all the forecasted horizons. As shown in Table 2, our *TOA & all satellite* model, which also avoids the use of telemetry, outperforms the previous model with errors of 19.9 % for the 1-h horizon, 36.8 % for the 6-h horizon, and a mean value for all forecasted horizons of 29.3 %. Although the numbers are never directly comparable because of the local climatology, the obtained results with *TOA & all satellite* are in the same range than those obtained in similar studies.

Table 2: Forecasting errors for the considered models with a 1-h granularity

		1h	2h	3h	4h	5h	6h
rRMSE (%)	Persistent	20.1	31.5	40.2	47.0	52.6	56.3
	TOA & all real	14.4	21.0	25.2	28.6	32.4	35.0
	Clear-sky & all real	14.3	20.8	24.9	28.4	31.7	33.8
	TOA & real labels	16.4	22.0	26.1	29.9	32.9	35.3
	TOA & all satellite	19.9	24.2	28.7	31.8	34.6	36.8
Skill (%)	TOA & all real	29.1	33.9	38.1	39.9	39.1	38.5
	Clear-sky & all real	29.8	34.5	38.8	40.2	40.4	40.6
	TOA & real labels	19.3	30.9	35.8	37.1	38.1	38
	TOA & all satellite	2.1	24.0	29.4	33.1	34.9	35.4

4.5. Comparison with the European Centre for Medium-Range Weather Forecasts (ECMWF) prediction

Table 3 compares the performance of the ECMWF forecast versus the proposed model with different inputs. Note that the




Figure 8: Relative RMSE with increasing horizons, 1-hour granularity.

Table 3: Mean forecasting error for 6 h in the forecasting horizon.

	MAE	rMAE	RMSE	rRMSE
ECMWF	67.6	17.4	97.2	24.9
TOA & all real	57.3	14.7	80.4	20.7
TOA & all satellite	63.7	16.4	89.6	23.0

proposed model outperforms the ECMWF model, even without irradiance measurements. Taking the latter as the reference model instead of the *persistent*, the forecast skill of our model is 17.3 % and 7.8 % for *TOA & all real* and *TOA & all satellite*, respectively.

5. Conclusion

This work introduces a DNN for the intra-day forecasting of solar irradiance that could help to improve the controllability of PV power plants with or without energy storage. The proposed DNN combines different data sources as well as different approaches such as cloud motion and machine learning to forecast the coming climatological conditions evolution. Furthermore, a study of the accuracy of the model considering the availability of the data sources has been conducted. Mainly, four case studies have been considered. First, only satellite data are used; second, actual data from the location are available only during the training of the DNN, but not later; third, the possibility of using actual data gathered in the target location as a feedback to the forecasting system is considered; last, the weather information for the future, in the form of a clear-sky irradiance forecast, is used.

As shown in the results, the DNN is capable of providing an accurate forecast of the solar irradiance, and it outperforms the *persistent* algorithm in all cases. Moreover, the proposed model outperforms the available NWP forecasts for the same

forecasting horizon and has a state-of-the-art performance even when no actual data are available.

Regarding the data input, the system benefits from the use of actual data in very short-term predictions as it helps to reduce the error, whereas the weather information proves to be useful in longer term predictions.

In future work, new data sources such as NWP forecasts for different variables such as wind speed, wind direction, and cloud cover as well as contamination forecasts or even multiple satellite images will be considered to add this information to the DNN. Furthermore, different structures that exploit the temporal part of the problem, specially LSTM layers, will be studied.

Acknowledgment

The authors would like to thank the financial support provided by the Universitat Jaume I from Castelló (Spain), the Generalitat Valenciana (GV), and the European Social Fund (ESF). This work was developed within the context of the projects with codes UJI-B2017-26 and grant number ACIF/2019/106.

References

- Alonso-Montesinos, J., & Batlles, F. (2015). Solar radiation forecasting in the short- and medium-term under all sky conditions. *Energy*, 83, 387–393.
- Bird, R. E., Riordan, C., Bird, R. E., & Riordan, C. (1986). Simple Solar Spectral Model for Direct and Diffuse Irradiance on Horizontal and Tilted Planes at the Earth’s Surface for Cloudless Atmospheres. *Journal of Climate and Applied Meteorology*, 25, 87–97.
- Blaga, R., Sabadus, A., Stefu, N., Dughir, C., Paulescu, M., & Badescu, V. (2019). A current perspective on the accuracy of incoming solar energy forecasting. *Progress in Energy and Combustion Science*, 70, 119–144.
- Brenna, M., Foadelli, F., Longo, M., & Zaninelli, D. (2016). Energy Storage Control for Dispatching Photovoltaic Power. *IEEE Transactions on Smart Grid*, 9, 2419–2428.
- Dong, Z., Yang, D., Reindl, T., & Walsh, W. M. (2014). Satellite image analysis and a hybrid ESSS/ANN model to forecast solar irradiance in the tropics. *Energy Conversion and Management*, 79, 66–73.
- ECMWF (2016). Copernicus Service Product Portfolio, .
- Greuell, W., Meirink, J. F., & Wang, P. (2013). Retrieval and validation of global, direct, and diffuse irradiance derived from SEVIRI satellite observations. *Journal of Geophysical Research: Atmospheres*, 118, 2340–2361.
- Gutierrez-Corea, F.-V., Manso-Callejo, M.-A., Moreno-Regidor, M.-P., & Manrique-Sancho, M.-T. (2016). Forecasting short-term solar irradiance based on artificial neural networks and data from neighboring meteorological stations. *Solar Energy*, 134, 119–131.
- Hanna, R., Kleissl, J., Nottrott, A., & Ferry, M. (2014). Energy dispatch schedule optimization for demand charge reduction using a photovoltaic-battery storage system with solar forecasting. *Solar Energy*, 103, 269–287.
- Jaeger-Waldau, A. (2018). *PV Status Report 2018*. Technical Report JRC Science for policy report.
- Jossen, A., Garche, J., & Sauer, D. U. (2004). Operation conditions of batteries in PV applications. *Solar Energy*, 76, 759–769.
- Katiraei, F., & Aguero, J. (2011). Solar PV Integration Challenges. *IEEE Power and Energy Magazine*, 9, 62–71.
- Lago, J., De Brabandere, K., De Ridder, F., & De Schutter, B. (2018). Short-term forecasting of solar irradiance without local telemetry: A generalized model using satellite data. *Solar Energy*, 173, 566–577.
- Larson, D. P., & Coimbra, C. F. M. (2018). Direct Power Output Forecasts From Remote Sensing Image Processing. *Journal of Solar Energy Engineering*, 140.
- Lauret, P., Voyant, C., Soubdhan, T., David, M., & Poggi, P. (2015). A benchmarking of machine learning techniques for solar radiation forecasting in an insular context. *Solar Energy*, 112, 446–457.

Marcos, J., Marroyo, L., Lorenzo, E., & García, M. (2012). Smoothing of PV power fluctuations by geographical dispersion. *Progress in Photovoltaics: Research and Applications*, 20, 226–237.

Mazorra Aguiar, L., Pereira, B., David, M., Díaz, F., & Lauret, P. (2015). Use of satellite data to improve solar radiation forecasting with Bayesian Artificial Neural Networks. *Solar Energy*, 122, 1309–1324.

Qazi, A., Fayaz, H., Wadi, A., Raj, R. G., Rahim, N., & Khan, W. A. (2015). The artificial neural network for solar radiation prediction and designing solar systems: a systematic literature review. *Journal of Cleaner Production*, 104, 1–12.

Quesada-Ruiz, S., Linares-Rodríguez, A., Ruiz-Arias, J., Pozo-Vázquez, D., & Tovar-Pescador, J. (2015). An advanced ANN-based method to estimate hourly solar radiation from multi-spectral MSG imagery. *Solar Energy*, 115, 494–504.

Ren, Y., Suganthan, P., & Srikanth, N. (2015). Ensemble methods for wind and solar power forecasting—A state-of-the-art review. *Renewable and Sustainable Energy Reviews*, 50, 82–91.

Segarra-Tamarit, J., Perez, E., Beltran, H., & Perez, J. (2020). Global horizontal irradiance forecasting code and data. Mendeley Data.

Tuohy, A., Zack, J., Haupt, S. E., Sharp, J., Ahlstrom, M., Dise, S., Gritmit, E., Mohrlén, C., Lange, M., Casado, M. G., Black, J., Marquis, M., & Collier, C. (2015). Solar Forecasting: Methods, Challenges, and Performance. *IEEE Power and Energy Magazine*, 13, 50–59.

Voyant, C., Notton, G., Kalogirou, S., Nivet, M.-L., Paoli, C., Motte, F., & Fouilloy, A. (2017). Machine learning methods for solar radiation forecasting: A review. *Renewable Energy*, 105, 569–582.

Wald, L., Hoyer-Klick, C., Lefèvre, M., & Schroedter-Homscheidt, M. (2015). *USER'S GUIDE to the MACC-RAD Services on solar energy radiation resources March 2015*. Technical Report.

Yang, D. (2019). A guideline to solar forecasting research practice: Reproducible, operational, probabilistic or physically-based, ensemble, and skill (ROPES). *Journal of Renewable and Sustainable Energy*, 11, 022701.

Yang, D., Kleissl, J., Gueymard, C. A., Pedro, H. T., & Coimbra, C. F. (2018). History and trends in solar irradiance and PV power forecasting: A preliminary assessment and review using text mining. *Solar Energy*, 168, 60–101.

Yang, D., Wu, E., & Kleissl, J. (2019). Operational solar forecasting for the real-time market. *International Journal of Forecasting*, 35, 1499–1519.

Zhang, J., Zhao, L., Deng, S., Xu, W., & Zhang, Y. (2017). A critical review of the models used to estimate solar radiation. *Renewable and Sustainable Energy Reviews*, 70, 314–329.

Zhao, X., Wei, H., Wang, H., Zhu, T., & Zhang, K. (2019). 3D-CNN-based feature extraction of ground-based cloud images for direct normal irradiance prediction. *Solar Energy*, 181, 510–518.

A. Supplementary error metrics

As discussed in Section 3.1 supplementary error metrics (MAE, rMAE and RMSE) are introduced here in Tables 4 and 5 for the sake of completeness and easier comparison with other works.

Table 4: Forecasting errors for the considered models with a 1-h granularity

		1h	2h	3h	4h	5h	6h
MAE (W/m ²)	Persistent	44.6	79.0	105.9	124.7	138.4	147.6
	TOA & all real	33.9	51.2	61.8	71.0	80.6	89.9
	Clear-sky & all real	33.8	50.3	61.1	70.4	80.0	86.7
	TOA & real labels	40.7	54.2	64.9	75.8	83.8	92.4
	TOA & all satellite	55.6	64.3	74.3	79.9	86.5	91.7
rMAE (%)	Persistent	11.9	21.1	28.2	33.2	36.9	39.3
	TOA & all real	9.1	13.8	16.6	19.1	21.7	24.3
	Clear-sky & all real	9.1	13.5	16.5	19.0	21.6	23.4
	TOA & real labels	11.0	14.6	17.5	20.4	22.6	24.9
	TOA & all satellite	15.0	17.3	20.0	21.5	23.3	24.7
RMSE (W/m ²)	Persistent	75.5	118.1	150.9	176.3	197.3	211.1
	TOA & all real	53.5	78.1	93.4	106.0	120.1	129.9
	Clear-sky & all real	53.0	77.3	92.4	105.4	117.6	125.4
	TOA & real labels	60.9	81.6	96.9	110.9	122.1	130.9
	TOA & all satellite	73.9	89.7	106.5	117.9	128.5	136.3

Table 5: Forecasting errors for the considered models with a 15-min granularity

		15min	30min	...	1h	...	2h	...	3h	...	4h	...	5h	...	5h45min	6h
MAE (W/m^2)	Persistent	35.7	49.0	-	65.6	-	98.4	-	123.3	-	140.6	-	152.0	-	158.7	159.9
	TOA & all real	33.9	42.6	-	52.6	-	67.9	-	77.3	-	86.5	-	96.5	-	103.2	108.0
	Clear-sky & all real	34.7	42.4	-	53.5	-	65.2	-	77.0	-	86.9	-	96.3	-	101.6	101.4
	TOA & real labels	45.1	49.3	-	58.1	-	70.5	-	81.1	-	92.2	-	99.3	-	106.7	108.7
	TOA & all satellite	63.0	63.2	-	69.1	-	78.6	-	86.0	-	92.0	-	99.7	-	103.4	105.1
rMAE (%)	Persistent	8.9	12.3	-	16.4	-	24.7	-	30.9	-	35.2	-	38.1	-	39.8	40.1
	TOA & all real	8.5	10.7	-	13.2	-	17.1	-	19.5	-	21.8	-	24.3	-	26.0	27.2
	Clear-sky & all real	8.7	10.7	-	13.5	-	16.4	-	19.4	-	21.9	-	24.3	-	25.6	25.6
	TOA & real labels	11.3	12.4	-	14.6	-	17.8	-	20.4	-	23.2	-	25.0	-	26.9	27.4
	TOA & all satellite	15.9	15.9	-	17.4	-	19.8	-	21.6	-	23.2	-	25.1	-	26.1	26.5
RMSE (W/m^2)	Persistent	71.8	89.8	-	107.9	-	144.8	-	175.4	-	197.9	-	216.4	-	226.1	227.9
	TOA & all real	61.8	72.3	-	84.4	-	103.5	-	115.2	-	127.0	-	139.0	-	145.1	149.6
	Clear-sky & all real	61.5	72.2	-	84.7	-	99.9	-	113.9	-	125.4	-	136.7	-	142.5	142.6
	TOA & real labels	70.3	77.3	-	89.0	-	104.9	-	118.8	-	131.4	-	141.1	-	147.1	148.8
	TOA & all satellite	86.8	88.6	-	97.4	-	112.9	-	126.4	-	136.8	-	146.8	-	151.8	154.0