

Meteorological data and Sky Images meets Neural Models for Photovoltaic Power Forecasting

Inés Montoya-Espinagosa¹ and Antonio Agudo², *Member, IEEE*

Abstract—Due to the rise in the use of renewable energies as an alternative to traditional ones, and especially solar energy, there is increasing interest in studying how to address photovoltaic forecasting in the face of the challenge of variability in photovoltaic energy production, using different methodologies. This work develops a hybrid approach for short and long-term forecasting based on two studies with the same purpose. A multimodal approach that combines images of the sky and photovoltaic energy history with meteorological data is proposed. The main goal is to improve the accuracy of ramp event prediction, increase the robustness of forecasts in cloudy conditions, and extend capabilities beyond nowcasting, to support more efficient operation of the power grid and better management of solar variability. Deep neural models are used for both nowcasting and forecasting solutions, incorporating individual and multiple meteorological variables, as well as an analytical solar position. The results demonstrate that the inclusion of meteorological data, particularly the surface long-wave, radiation downwards, and the combination of wind and solar position, significantly improves current predictions in both nowcasting and forecasting tasks, especially on cloudy days. This study highlights the importance of integrating diverse data sources to improve the reliability and interpretability of solar energy prediction models.

I. INTRODUCTION

Photovoltaic (PV) solar energy has become a crucial element in the global transition to renewable energy [1], and is projected to become the largest renewable source by 2030 [2], [3], [4]. However, the variability of solar energy, mainly due to the dynamic nature of clouds, causes rapid fluctuations in irradiation resulting in an unstable energy production [5], [6], which negatively impacts grid stability [7], [8], [9]. This unpredictability makes intra-hour solar forecasting and short-term photovoltaic forecasting particularly difficult, yet essential for maintaining grid stability and ensure safe system operation [10], [11], [12], [13], [14], [15]. To address these challenges, different studies aim to improve the accuracy of solar forecasts using sky image-based prediction models [16], [17], some physical information such as solar irradiance and temperature [18] or perceptible water [19] were used.

An approach to improve forecast accuracy involves using sky image-based prediction models, which leverage visual information to anticipate PV fluctuations more accurately than purely numerical and statistical models [20], [13].

This work has been supported by the project GreenVAR of the Fundación Ramón Areces.

¹Inés Montoya-Espinagosa is with the Universitat Pompeu Fabra, Barcelona, Spain.

²Antonio Agudo is with the Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain.

For instance, [21] developed a two-stream network with PV energy-guided attention, achieving high accuracy but struggling with limited data and a lack of interpretability. Similarly, [20] used a deep-learning model for irradiance nowcasting and integrated it into a hybrid system with physics-based models and smart persistence, improving persistence models, but still facing difficulties predicting ramp events (sudden changes in power). Other studies have explored hybrid deep-learning architectures. For example, [22] developed models using static and dynamic sky images, finding that their hybrid model achieved higher performance but generated blurry cloud boundary predictions. [23] used cloud coverage rates for higher accuracy but the method fails in rainy conditions and encounters limitations in RGB-based coverage calculations. To address fog-affected scenarios, [24] introduced the aerosol scattering coefficient as an auxiliary feature, achieving notable improvements when combined with other algorithms based on support vector machines, long short-term memory (LSTM) models, and XGBoost. [25] used cloud motion vectors for high accuracy but resulted in an overfitted model heavily relies on accurate cloud tracking. Collectively, these studies underscore the need for robust, multi-modal approaches that can adapt to diverse weather conditions. Meanwhile, [26] proposed a hybrid mapping model combining Convolutional Neural Networks (CNN) and LSTM, which outperformed standalone models in various metrics but required further improvement in generalizing to varied cloud dynamics. [17] used CNNs with the Sky Images and Photovoltaic Power Dataset (SKIPP'D) for PV prediction and in [27] incorporated image-based estimations such as cloud and sun segmentation or sun position to increase the robustness of the method. However, these approaches faced key challenges, including the inability to learn all important patterns, reduced reliability during cloudy weather with rapid power changes, as well as issues with image brightness and cloud identification. The increasing use of solar energy into power grid has revealed a critical problem: solar energy production varies a lot due to the cloud's motion and changing weather. While artificial-intelligence methods have helped with short-term solar predictions, current systems still have two key limitations: (1) efficiently predict fluctuations in solar irradiance (called ramp events), and (2) perform accurate forecasts for very short timeframe. These challenges have a direct impact on the grid stability and operational efficiency, since they need to perfectly match electricity supply with demand at all times.

To address the previous gaps, this work proposes a multi-modal hybrid methodology that combines ground-based sky

images from SKIPP'D [17] with physical meteorological data from the ERA5 reanalysis [28]. The objective is to enhance the prediction of ramp events, improve forecast robustness, extend capabilities beyond nowcasting, and provide a physical-aware interpretability to support more efficient and reliable power grid operation.

II. METEOROLOGICAL AND VISUAL INFORMATION

We now introduce the information we propose to use in the prediction of PV power. To this end, we consider the dataset SKIPP'D [17] for visual information and ERA5 [29] for meteorological one.

A. Visual information

The SKIPP'D dataset [17] comprises ground-based sky images captured at a frequency of 20 Hz and a high spatial resolution of $2,048 \times 2,048$ pixels using a 360-degree fisheye camera. The recordings span the period from March 2017 to December 2019 and cover the time interval between 06:00 AM and 08:00 PM each day. In addition to the imagery, the dataset includes PV power generation data recorded at 1-minute intervals. Following the work described in [17], [27], we utilize 363,375 sky image samples with a temporal resolution of one minute and a downsampled spatial resolution of 64×64 pixels (some instances are displayed in Fig. 2), along with their corresponding PV output observations, and employing a 96%/4% split for training and testing, respectively. The dataset contains 20 selected days: 10 sunny and 10 cloudy.

From that data, we derived an additional dataset for forecasting purposes, referred to as the Forecast dataset. This contains 137,489 samples, approximately half the size of the original, by applying a two-minute sampling interval each day. In this particular case, it has been employed a 94.25%/5.75% split for training and testing, respectively. It is worth noting that in this case sequences of 16 consecutive images (covering the past 15 minutes) and the corresponding 16 PV power values are used.

B. Meteorological information

The ERA5 [29] reanalysis dataset from ECMWF [28] provides a globally consistent, hourly record of meteorological variables from 1940 to present. Data are provided on a 0.25-degree grid (approximately 31 km in resolution). Next, we introduce some variables that will be later analyzed in PV power prediction:

1) *Data Coverage and Extraction*: Meteorological data are extracted from four grid points (NW, NE, SW, SE) surrounding the camera location at Stanford University (37.427°N , -122.174°W), as shown in Fig. 1. The spatial coverage of this area is displayed in the figure.

2) *Meteorological variables*: Ten key variables have been extracted, categorized as (a summary of meteorological data is displayed in Table I):

- **Instantaneous analysis variables**: Total cloud cover (tcc), surface pressure (sp), and wind components (Wind100 which is composed by $100u$ and $100v$).




Fig. 1: **Latitude and longitude of the four ERA5 [29] grid points** surrounding the camera location in Stanford University. These points, spaced at 0.25° resolution (NW, NE, SW, SE), define the area from which ERA5 [29] variables are extracted. The red square indicates the real camera location for image acquisition.

- **Forecast non-accumulative variable**: Wind gust ($i10fg$).
- **Forecast accumulative variables**: Radiation fluxes: ($ssrd$, $strd$, str , tsr , $fdir$).

3) *Meteorological data preprocessing*: Three different preprocessing techniques have been applied based on variable type:

- **Instantaneous analysis variables**: Converted from UTC to local time, extracted from the four grid points (NW, NE, SW, SE), aligned to image timestamps using forward/backward fill, and normalized.
- **Forecast non-accumulative variable**: Real timestamps reconstructed using $time + step$, then processed similarly to analysis variables.
- **Forecast accumulative variables**: Converted from cumulative to hourly rates, then flattened, assigned real timestamps, and interpolated to 1-minute resolution.

Mean and standard deviation are computed on the training and validation split, forcing standard deviation to 1.0 to avoid division by zero. This normalization ensures that all variables contribute equally during model training, regardless of their original units and/or magnitudes.

C. Sun Position

Finally, we also propose the sun position as an additional source of information to improve model prediction by helping identify lighting patterns and atmospheric conditions linked to the solar cycle.

To this end, we propose an analytical way to infer the solar location by exploiting two fundamental solar angles:

- Solar azimuth (α): horizontal angle from true north (0° – 360°).
- Solar elevation (ϵ): angle between the horizon and the sun (0° – 90°).

These parameters were obtained via `pvlib` [30], a Python package specifically designed for solar energy system modeling. For the fisheye camera at Stanford University (latitude 37.427° , longitude -122.174° , oriented 194° from true north),

TABLE I: ERA5 [29] variables (**fc** = forecast; **an** = analysis) explanation. The table shows, from left to right, a summary of the official parameter name of the meteorological variables, the full name, the source (if it is an analysis or forecast variable), whether the data has or not a step, the units, and the variable type (whether it is an accumulative variable in time or not).

Variable	Full name	Source	Step	Units	Variable type
tcc	Total Cloud Cover	an	No	Fraction (0-1)	Instantaneous
i10fg	10m Wind Gust	fc	Yes	m/s	Instantaneous
100u	U component of wind at 100m	an	No	m/s	Instantaneous
100v	V component of wind at 100m	an	No	m/s	Instantaneous
sp	Surface Pressure	an	No	Pa	Instantaneous
strd	Surface Thermal Radiation Downwards	fc	Yes	J/m ²	Accumulative
ssrd	Surface Solar Radiation Downwards	fc	Yes	J/m ²	Accumulative
str	Surface Net Thermal Radiation	fc	Yes	J/m ²	Accumulative
tsr	Top Net Solar Radiation	fc	Yes	J/m ²	Accumulative
fsr	Surface Solar Direct Radiation	fc	Yes	J/m ²	Accumulative

the solar azimuth is first converted to a relative azimuth with respect to the camera’s optical axis as:

$$\alpha_{\text{rel}} = (\alpha_{\text{sun}} - \alpha_{\text{cam}}) \text{ with } \text{mod } 360^\circ, \quad (1)$$

where α_{rel} is the corrected relative azimuth, α_{sun} is the solar azimuth and $\alpha_{\text{cam}} = 194^\circ$ indicates the camera’s orientation. The relative azimuth is then converted to radians as $\theta = \text{deg2rad}(\alpha_{\text{rel}})$.

The radial distance from the image center is calculated using a normalized zenith angle:

$$\theta_{\text{norm}} = \frac{90^\circ - \epsilon}{90^\circ}, \quad (2)$$

that it is used to infer the radial distance r as:

$$r = R \cdot (\theta_{\text{norm}})^\gamma, \quad (3)$$

where R is the maximum image radius (32 pixels for 64×64 images) and $\gamma = 1.2$ denotes a non-linear correction factor.

A vertical offset correction is also applied to account for camera tilt such that:

$$\delta_y = y_0 + m(\epsilon - 45^\circ), \quad (4)$$

with $y_0 = -3.0$ and $m = 0.03$. The final (x, y) coordinates in the image plane (origin at upper-left, y -axis downward) are then inferred as:

$$x = C + r \cdot \sin(\theta), \quad (5)$$

$$y = C + r \cdot \cos(\theta) + \delta_y, \quad (6)$$

where C is the image center (e.g., 32 pixels for 64×64 images). In spite of not having a ground truth for evaluation, we provide some estimations of the resulting sun position in different scenarios in Fig. 2. As it can be seen, the results show some alignment trade-offs between center and corner positions, despite the empirical corrections applied throughout the transformation process. However, at least visually, the estimation seems correct and accurate to be included in our physical-aware PV power estimation model. In fact, this solution seems to be more robust than that obtained in [27], where an image processing approach was used for this task, especially when the sky is very cloudy.




Fig. 2: Analytical sun position estimation. Four sky input images in different conditions where our sun position estimation is displayed by means of a \times red mark. Best viewed in color.

III. NOWCAST AND FORECAST PV POWER ESTIMATION

In this paper, and following [17], [27], we present two neural models for nowcast and forecast PV power estimation. In the first case, our method retrieves an instantaneous power estimation from a single sky image. In addition to that, as this is a short-term prediction where high-accuracy information could be available, we may exploit real-time observations such as those coming from radars, satellites or weather stations; for rapidly evolving events like thunderstorms, heavy precipitation, or fog just to name a few. Particularly, we propose to exploit meteorological data to improve our neural model. Moreover, in this context the spatial scale is local (microscale and gamma mesoscale, few kilometers of resolution). Our model f_N can be represented as:

$$f_N : (I_i, M_i) \rightarrow P_i, \quad (7)$$

where I_i represents the current image, M_i the meteorological information and P_i the PV power estimation at the i -th time.

In the second case, the forecasting model, the PV power production of the last 16 minutes together with the corresponding sky images are considered for a 15-minute ahead prediction. Our model f_F can be written as:

$$f_F : (I, M, P)_{t_i-H:\delta:t_i} \rightarrow P_{t_i+T}, \quad (8)$$

where $H = 15$ minutes is the length of historical terms, $\delta = 1$ min the interval between historical terms, a 2 minute frequency sampling for the forecasting case, and $T = 15$ minutes represents the prediction horizon. Then, the main difference between both models is the use of information: nowcasting uses only instantaneous inputs and forecasting incorporates historical context sequences. In both cases the output power estimation is given in kilowatts (kW).

Fig. 3 shows the overall structure of the nowcast and forecast models, where the differences in the input data that feed into each architecture can be seen. Without loss of generality, both models are quite similar in terms of the main architecture we exploit. In particular, two blocks with convolutional layers with 3×3 filters with a stride of 1 and same-value padding, having 24 and 48 filters, respectively; as well as batch normalization and a rectified linear unit function for activation are considered. The last part of the block being a 2×2 max-pooling layer with a stride of 2. After that, flattening the tensors the data are sent to the third and last block, which consists of two fully connected layers with 1,024 neurons each and a dropout of 40% to enhance stability of the model. Before the FC layers, this is where the difference between the two models lies, as it is here




Fig. 3: Architectures of the neural model for both nowcast (left) and forecast (right) tasks. The diagram shows the different blocks, the input variables fed into the model, and the output shapes at each stage.

that the historical PV data are considered, along with the meteorological variables we propose to use, i.e., all variables are concatenated with flattened image features, introducing physical information directly into the latent feature space. To the best of our knowledge, this is the first time this type of information has been used in order to sort out this problem.

In particular, three sets of variables are considered to constrain both nowcast and forecast models:

- **Single meteorological variable.** 4-element vector representing values from NW, NE, SW, SE grid points.
- **Multiple meteorological variables.** Concatenated vector of different combinations of meteorological variables, see them in section II-B.3 (length = $4 \times$ number of variables).
- **Sun position.** Normalized coordinates (sun_x , sun_y), with shape 1D and two values.

Both models are trained using Adam optimizer [31] with a learning rate of 3×10^{-6} . In training, a mean squared error loss is minimized between the ground truth and the prediction.

IV. EXPERIMENTAL RESULTS

In this section, we present our experimental evaluation for both nowcast and forecast models as well as a comparison with state-of-the-art approaches [17], [27]. Following literature, model performance is evaluated using the Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) as:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{n=1}^N (y_n - \hat{y}_n)^2}, \quad (9)$$

$$\text{MAE} = \frac{1}{N} \sum_{n=1}^N |y_n - \hat{y}_n|, \quad (10)$$

where y_n and \hat{y}_n denote the n -th ground truth value and the prediction, respectively; and N the number of samples.

Table II reports our experimental results for both nowcast and forecast PV power models, including the state-of-the-art approaches SUNSET [17] and RSUNSET [27] as well as our own implementation of the SUNSET model denoted as MSUNSET. Using MSUNSET as a baseline where just the sky images are considered, different one-by-one meteorological variables (see section II-B.3) are exploited to finally be combined and concatenated with the sun position we propose. In general, it is worth pointing out the difference

in performance between sunny (0.73/2.33kW) and cloudy (3.32/2.49kW) days for nowcast/forecast tasks, respectively, being the cloudy scenario more challenging. Next, we analyze the results in depth.

Nowcast results. The variable tsr (top net solar radiation) improved predictions on cloudy days, contrary to intuitive expectations. This suggests that atmospheric radiation attenuation may serve as an indicator of cloud presence and density, and model would learn the interception of solar radiation before reaching the surface. Then, regarding the rest of the radiative variables (strd , ssrd , and str), performed better on sunny days, likely due to simplified radiative states under clear conditions.

Surface pressure (sp) showed limited improvement on cloudy days despite its theoretical relationship with wind patterns. In fact, $\text{i10fg} + \text{wind100}$ provided a consistent improvement in prediction. In particular, the combination $\text{i10fg} + \text{wind100} + \text{strd}$ significantly worsened the results despite the improvements in individual cases, indicating the difficulty of the model in capturing complex variable interactions. However, $\text{i10fg} + \text{wind100} + \text{sun position}$ achieves one of the most accurate solutions, even outperforming competing methods [17], [27], suggesting that the context provided by the solar position along with other effective predictor variables increases its effectiveness.

Forecast results. In forecasting, we were not able to exactly reproduce the results of the forecast models reported in [17], [27], despite we deduced the dataset ourselves by following the procedure in the original papers. As it can be seen, our baseline provides in terms of RMSE better results in cloudy days but worse in sunny ones, probably due to the final partitioning of the data. The model's specific improvement on cloudy days, despite clear days being inherently more predictable, suggests it has effectively learned to handle the complexity and variability characteristic of cloudy scenarios. Then, the improvement in RMSE but not MAE indicates that the model has successfully reduced extreme errors, as RMSE penalizes large errors while MAE treats all errors equally.

This time, we only consider the best combinations in nowcasting. On the one hand, when our model is considering the variables (i10fg , wind100) and sun position, the solution is good on average, but it is not capable of surpassing the base model, suggesting that MSUNSET adapts better to cloudy conditions without variable complexity. Contrary to expectations, combining solar position with wind variables enhances forecasting in cloudy rather than sunny conditions, as opposed to the nowcast model. This unexpected result suggests both nowcast and forecast models effectively learn meaningful patterns from wind-sun position interactions. On the other hand, the best solution on average is achieved by exploiting the variable strd , providing a greatest improvement in cloudy days, contrasting with the performance in the nowcast model, where it excels in sunny conditions. This difference suggests that forecast models leverages historical information that allowed to better use thermal patterns associated with cloud cover, while on clear days it may be less relevant or could even interfere

TABLE II: Nowcast and Forecast PV power predictions in sunny, cloudy and overall scenarios in terms of RMSE and MAE. Nowcast (top) and forecast (bottom) results as a function of several sources of information. The table also reports the results provided by SUNSET [17] and RSUNSET [27] models. * means RSUNSET + sun position + SAMPI results, the best solution in [27]. In bold are marked the best models.

Model	RMSE↓	MAE↓
	(sunny / cloudy / overall)	(sunny / cloudy / overall)
<i>Nowcast results</i>		
SUNSET [17]	0.804 / 3.335 / 2.428	0.657 / 2.337 / 1.499
RSUNSET* [27]	0.790 / 3.300 / 2.400	0.650 / 2.300 / 1.480
MSUNSET (Ours)	0.804 / 3.335 / 2.428	0.657 / 2.337 / 1.499
+tcc	0.783 / 3.343 / 2.430	0.643 / 2.345 / 1.496
+i10fg	0.764 / 3.305 / 2.401	0.636 / 2.302 / 1.471
+ssrd	0.768 / 3.355 / 2.436	0.640 / 2.360 / 1.502
+str	0.746 / 3.369 / 2.442	0.612 / 2.370 / 1.493
+tsr	0.828 / 3.332 / 2.430	0.678 / 2.327 / 1.504
+strd	0.732 / 3.316 / 2.404	0.600 / 2.319 / 1.461
+sp	0.847 / 3.333 / 2.434	0.708 / 2.336 / 1.524
+fdir	0.821 / 3.358 / 2.447	0.681 / 2.370 / 1.527
+Wind100	0.798 / 3.328 / 2.423	0.673 / 2.320 / 1.498
+sun position	0.791 / 3.372 / 2.452	0.645 / 2.374 / 1.512
+i10fg+Wind100+strd	0.864 / 3.283 / 2.403	0.756 / 2.266 / 1.513
+i10fg+Wind100	0.782 / 3.314 / 2.410	0.660 / 2.285 / 1.474
+i10fg+strd+Wind100+tcc	0.820 / 3.297 / 2.405	0.699 / 2.273 / 1.488
+i10fg+strd+str+Wind100	0.818 / 3.300 / 2.406	0.703 / 2.281 / 1.493
+i10fg+Wind100+sun position	0.799 / 3.297 / 2.401	0.689 / 2.276 / 1.484
+i10fg+Wind100+sun position+strd	0.826 / 3.302 / 2.409	0.707 / 2.274 / 1.493
<i>Forecast results</i>		
SUNSET [17]	0.610 / 4.270 / 3.030	0.500 / 2.950 / 1.710
RSUNSET* [27]	8.620 / 7.610 / 8.310	7.550 / 5.750 / 6.650
MSUNSET (Ours)	2.461 / 2.539 / 2.500	2.109 / 1.938 / 2.023
+i10fg	2.589 / 2.615 / 2.603	2.197 / 2.037 / 2.117
+strd	2.331 / 2.489 / 2.412	1.924 / 1.882 / 1.903
+i10fg+Wind100	2.831 / 2.669 / 2.751	2.402 / 2.056 / 2.229
+i10fg+Wind100+sun position	2.570 / 2.494 / 2.532	2.269 / 1.934 / 2.101
+i10fg+Wind100+strd	3.056 / 2.762 / 2.912	2.606 / 2.088 / 2.345

with the model’s pattern learning. In contrast, nowcast model, despite having one sky image available, benefit from `strd` as a stable clear-sky indicator.

On balance, it is concluded that the use of `strd` as well as the wind parameters (`i10fg`, `wind100`) and solar position (`sun position`) provides the most significant improvements in both nowcast and forecast tasks. A qualitative evaluation in ten test sunny and cloudy days for that combination is displayed in Fig. 4. As it can be seen, our algorithm provides a competitive solution in both cases, while helping to obtain a better interpretability of the solution.

V. CONCLUSION

In this paper we demonstrate the effectiveness of a hybrid multimodal approach for PV power forecasting that integrates sky images, historical power generation data, and meteorological variables through deep neural models. The results confirm that incorporating specific meteorological parameters, particularly surface long-wave (thermal) radiation downwards, wind, and solar position, enhances both short- and long-term prediction accuracy, with marked improvements specially under cloudy conditions. These findings underline the value of combining diverse data modalities to mitigate the inherent variability of solar energy production and to strengthen the robustness and interpretability of predictive models. Then, this approach contributes to more reliable grid operation and efficient management of renewable energy resources, paving the way for future developments in adaptive and data-driven solar forecasting systems.

REFERENCES

- [1] R. Pérez-Gonzalo, A. Espersen, and A. Agudo, “Robust wind turbine blade segmentation from RGB images in the wild,” in *ICIP*, 2023.
- [2] “Global Energy Review 2025 – Analysis - IEA — [iea.org](https://www.iea.org/reports/global-energy-review-2025),” <https://www.iea.org/reports/global-energy-review-2025>, 2025.
- [3] R. Pérez-Gonzalo, A. Espersen, and A. Agudo, “Generalized nested latent variable models for lossy coding applied to wind turbine scenarios,” in *ICIP*, 2024.
- [4] S. Singhal, R. Pérez-Gonzalo, A. Espersen, and A. Agudo, “Dual-space augmented intrinsic-lora for wind turbine segmentation,” in *ICASSP*, 2025.
- [5] A. K. Barnes, J. C. Balda, and J. K. Hayes, “Modelling PV clouding effects using a semi-markov process with application to energy storage,” *IFAC*, vol. 47, no. 3, pp. 9444–9449, 2014.
- [6] L. Cozzi, T. Gould, S. Bouckart, D. Crow, T. Kim, C. Mcglade, P. Olejarnik, B. Wanner, and D. Wetzel, “World energy outlook 2020,” *IEA*, vol. 2050, pp. 1–461, 2020.
- [7] P. Kang and J. Li, “Short-term photovoltaic power forecasting based on VMD-KPCA-LSTM,” in *MLCCIM*, 2024, pp. 372–386.
- [8] J. Kim, J. Obregon, H. Park, and J. Y. Jung, “Multi-step photovoltaic power forecasting using transformer and recurrent neural networks,” *RSER*, vol. 200, 2024.
- [9] R. Othman, K. G. Herbert-Berger, and T. J. Marlowe, “An LSTM model for sustainable microgrid energy resources in myanmar,” in *CAI*, 2025.
- [10] E. Lorenz, J. Kühnert, and D. Heinemann, *Overview of Irradiance and Photovoltaic Power Prediction*, 2014.
- [11] F. Lin, Y. Zhang, and J. Wang, “Recent advances in intra-hour solar forecasting: A review of ground-based sky image methods,” *IJF*, vol. 39, pp. 244–265, 2023.
- [12] Y. Zahraoui, T. Korotko, S. Mekhilef, and A. Rosin, “ANN-LSTM based tool for photovoltaic power forecasting,” in *SGRE*, 2024.
- [13] X. Zhang, Z. Zhen, Y. Sun, F. Wang, Y. Zhang, H. Ren, H. Ma, and W. Zhang, “Prediction interval estimation and deterministic forecasting model using ground-based sky image,” *TIA*, vol. 59, pp. 2210–2224, 2023.

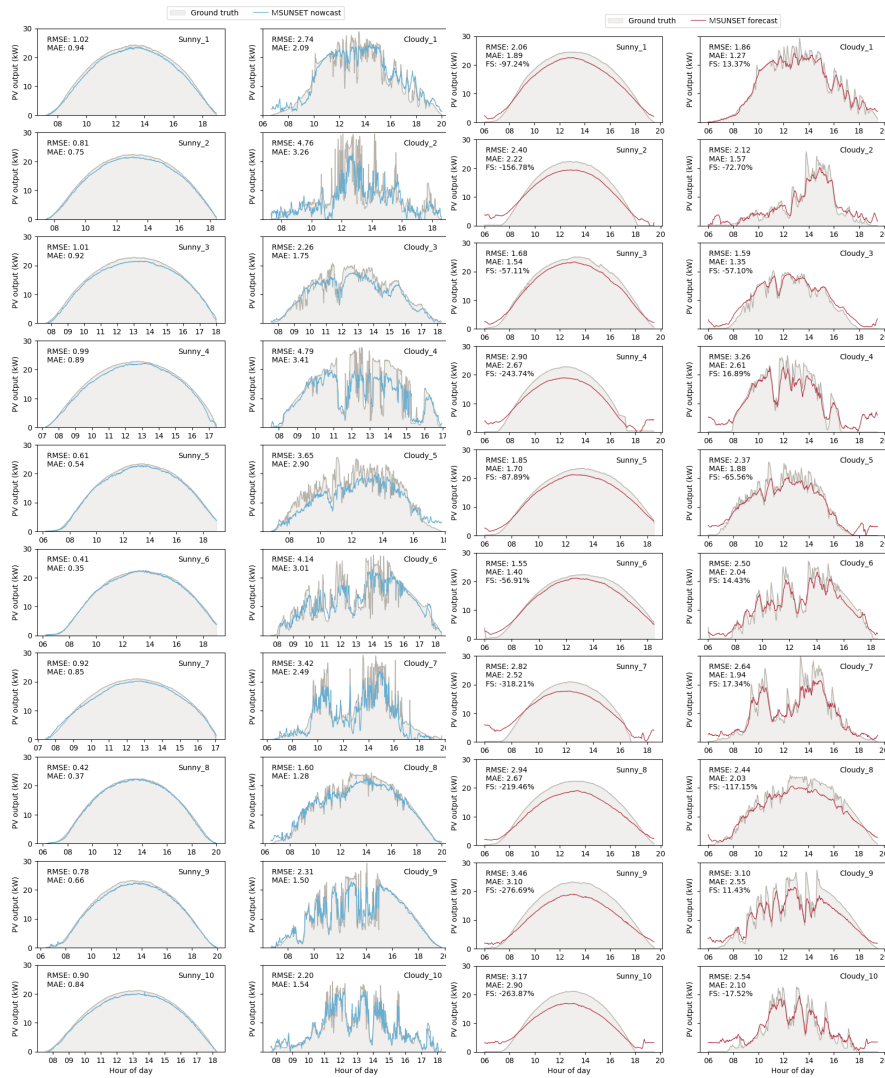


Fig. 4: Qualitative comparison on nowcast and forecast estimation for sunny and cloudy days. In both cases, we report our estimation MSUNSET + i10fg + wind100 + sun position. Grey area corresponds to the ground truth; blue and red lines represent the result of our model prediction for nowcast and forecast, respectively.

- [14] Y. Sun, V. Venugopal, and A. Brandt, "Short-term solar power forecast with deep learning: Exploring optimal input and output configuration," *Solar Energy*, vol. 188, pp. 730–741, 2019.
- [15] A. Mokri and V. Gadepally, "The efficiency frontier: Optimal combinations of energy-saving strategies in data centers," in *CAI*, 2025.
- [16] C. Feng, D. Yang, B. Hodge, and J. Zhang, "Opensolar: Promoting the openness and accessibility of diverse public solar datasets," *Solar Energy*, vol. 188, pp. 1369–1379, 2019.
- [17] Y. Nie, X. Li, A. Scott, Y. Sun, V. Venugopal, and A. Brandt, "SKIPP'D: a sky images and photovoltaic power generation dataset for short-term solar forecasting," *Solar Energy*, vol. 255, pp. 171–179, 2023.
- [18] U. Das, K. Tey, M. Seyedmahmoudian, M. Idna Idris, S. Mekhilef, B. Horan, and A. Stojcevski, "SVR-based model to forecast PV power generation under different weather conditions," *Energies*, vol. 10, no. 7, p. 876, 2017.
- [19] R. Ahmed, V. Sreeram, Y. Mishra, and M. Arif, "A review and evaluation of the state-of-the-art in PV solar power forecasting: Techniques and optimization," *RSER*, vol. 124, 2020.
- [20] Y. Fabel, B. Nouri, S. Wilbert, N. Blum, D. Schnaus, R. Triebel, L. F. Zarzalejo, E. Ugedo, J. Kowalski, and R. Pitz-Paal, "Combining deep learning and physical models: A benchmark study on all-sky imager-based solar nowcasting systems," *Solar RRL*, vol. 8, 2 2024.
- [21] H. Zang, D. Chen, J. Liu, L. Cheng, G. Sun, and Z. Wei, "Improving ultra-short-term photovoltaic power forecasting using a novel sky-image-based framework considering spatial-temporal feature interaction," *Energy*, vol. 293, 2024.
- [22] W. Kong, Y. Jia, Z. Y. Dong, K. Meng, and S. Chai, "Hybrid approaches based on deep whole-sky-image learning to photovoltaic generation forecasting," *Applied Energy*, vol. 280, 12 2020.
- [23] W. C. Kuo, C. H. Chen, S. Y. Chen, and C. C. Wang, "Deep learning neural networks for short-term pv power forecasting via sky image method," *Energies*, vol. 15, 7 2022.
- [24] Z. Lu, W. Chen, Q. Yan, X. Li, and B. Nie, "Photovoltaic power forecasting approach based on ground-based cloud images in hazy weather," *Sustainability (Switzerland)*, vol. 15, 12 2023.
- [25] K. Hu, L. Wang, W. Li, S. Cao, and Y. Shen, "Forecasting of solar radiation in photovoltaic power station based on ground-based cloud images and bp neural network," *IETGTD*, vol. 16, pp. 333–350, 2022.
- [26] Z. Zhen, J. Liu, Z. Zhang, F. Wang, H. Chai, Y. Yu, X. Lu, T. Wang, and Y. Lin, "Deep learning based surface irradiance mapping model for solar PV power forecasting using sky image," *TIA*, vol. 56, pp. 3385–3396, 2020.
- [27] A. Berresheim and A. Agudo, "Photovoltaic power forecasting using sky images and sun motion," in *ICASSP*, 2024.
- [28] "ECMWF Reanalysis v5," <https://www.ecmwf.int/en/forecasts/dataset/ecmwf-reanalysis-v5>, 2025.
- [29] C3S, "ERA5 hourly data on single levels from 1940 to present," 2018.
- [30] W. F. Holmgren, C. W. Hansen, and M. A. Mikofski, "Pvlib python: A python package for modeling solar energy systems," *J. Open Source Softw.*, vol. 3, no. 29, p. 884, 2018.
- [31] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *arXiv preprint arXiv:1412.6980*, 2014.