



HAL
open science

Improving cross-site generalisability of vision-based solar forecasting models with physics-informed transfer learning

Quentin Paletta, Yuhao Nie, Yves-Marie Saint-Drenan, Bertrand Le Saux

► To cite this version:

Quentin Paletta, Yuhao Nie, Yves-Marie Saint-Drenan, Bertrand Le Saux. Improving cross-site generalisability of vision-based solar forecasting models with physics-informed transfer learning. 2024. hal-04595797

HAL Id: hal-04595797

<https://hal.science/hal-04595797v1>

Preprint submitted on 31 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Improving cross-site generalisability of vision-based solar forecasting models with physics-informed transfer learning

Quentin Paletta^{a,b,c,*}, Yuhao Nie^{d,e}, Yves-Marie Saint-Drenan^f, Bertrand Le Saux^b

^a*Department of Engineering, University of Cambridge, UK*

^b*ESA Φ -lab, European Space Agency - ESRIN, Italy*

^c*ESA Climate Office, European Space Agency - ECSAT, UK*

^d*Department of Energy Science and Engineering, Stanford University, United States*

^e*Institute for Data, Systems, and Society, Massachusetts Institute of Technology, United States*

^f*MINES ParisTech, PSL Research University, O.I.E. Centre Observation, Impacts, Energy, France*

Abstract

Forecasting solar energy from cloud cover observations is crucial to truly anticipate future changes in power supply. On an intra-hour timescale, ground-level sky cameras located near a solar site offer the most valuable source of information on incoming clouds. In the literature, the analysis of these hyperlocal cloud cover observations for solar modelling is increasingly performed by deep learning algorithms trained and tested on years' worth of local data. However, this approach is not suitable for industrial applications since solar energy producers cannot wait for years of local data collection to start generating reliable solar forecasts. However, they might own relevant multi-location data collected from other solar sites over time. This study thus explores the capability of such algorithms to generalise beyond their training location in two data scarce conditions: zero-shot learning (i.e. direct application of a trained model to a new location without local fine-tuning) and few-shot learning (i.e. calibration of a pre-trained model based on very limited local data such as a day of observations). Zero-shot learning results show that using local clear-sky models to normalise output variables (e.g. solar irradiance or solar energy production values) facilitates cross-dataset transfer learning. Compared to previous methods, the resulting forecast skill increases by close to 25% in cloudy conditions and by more than 700% in clear-sky conditions. An additional gain is observed when local data collected in overcast weather conditions are used for model calibration via few-shot learning. The corresponding neural networks trained in data scarce conditions achieve comparable performance to expert local models based on years of training data. These promising results shed light on the potential of large-scale and multi-location sky image datasets to improve the generalisation skills of solar forecasting algorithms.

Keywords:

Solar Energy Forecasting, Computer Vision, Generalisability, Transfer learning, Zero-shot Learning, Few-shot Learning

*E-mail address: quentin.paletta@esa.int (Quentin Paletta)

1. Introduction

The current energy transition towards an increasing consumption of electricity together with higher portions of low carbon energy sources requires a deep transformation of our power systems [1, 2]. In particular, the variable nature of renewables such as wind and solar poses major challenges for their future large-scale integration in the energy mix [3, 4, 5]. To address this difficulty, several technical solutions have been explored including energy storage [6], spinning reserves [7] and demand flexibility [8]. To facilitate these methods, predicting the future energy demand and supply is crucial [9, 10, 11, 12].

In solar energy, the dominant source of uncertainty directly arises from diverse atmospheric factors including clouds and aerosols. If not anticipated on time, the resulting solar power fluctuations can have critical consequences on the security of energy supply, hence increasing the risk of power outage or market price volatility. For this reason, solar energy forecasting is expected to be predominantly tackled via the modelling of the atmosphere at different spatio-temporal scales [13]. As such, improving solar power predictions would benefit various activities including energy storage — pumped, battery, thermal and mechanical energy storage — and spinning reserve optimisation, energy trading, frequency setting and energy dispatch [4, 14, 15, 16, 17].

1.1. Solar forecasting with computer vision

On a short-term time scale from a few minutes to several hours ahead, computer vision algorithms offer a significant opportunity to model the cloud cover dynamics from Earth observations via remote and local sensing such as satellite or ground-based sky imagery [18, 19]. Having the advantage of processing data substantially faster than Numerical Weather Prediction (NWP) models, these methods based on physics or machine learning are able to integrate the most recent measurements to ameliorate short-term predictions up to several hours ahead [20]. In addition, NWP are based on physical model simulations, which have both uncertainties and limitations. In comparison, methods using earth observations benefit from direct measurements of the cloud cover state, and thus tend to be more accurate.

Recent advances in the field of computer-vision-aided solar forecasting have employed neural networks for image or video analysis [21]. In particular, convolutional neural networks, recurrent neural networks, and transformers have been trained to extract relevant spatio-temporal features from cloud cover images for solar modelling [22, 23, 24]. As such, deep learning models offer a significant potential to improve solar forecasts by utilising atmospheric observations from satellites [25], all-sky imagers (ASIs) [26] or a combination of both [27].

From that perspective, ASIs equipped with hemispherical sky cameras offer high spatio-temporal and hyper-local information on the cloud cover for solar power or irradiance map modelling. Following a decade development of physics-based algorithms based on cloud detection, tracking and radiative transfer modelling [28, 29, 30, 31], a wide range of data-driven methods have been proposed in recent years [32, 33, 34, 35, 36]. In particular, Zhang et al. [22] first benchmarked the application of deep learning and in particular recurrent neural networks to sky images for solar irradiance forecasting. Sun et al. [37] investigated diverse input and output configurations for PV power forecasting using a convolutional neural network. In this work, the variables of interests are the sampling frequency, as well as the number, interval and pacing of lag terms. To combine

sky image processing with auxiliary data, Siddiqui et al. [38] developed a neural network design with two parallel branches whose outputs are integrated in a single regression head. Paletta et al. [39] highlighted, in a benchmark study composed of common deep learning architectures, a consistent lag of data-driven predictions relative to the ground truth. Kong et al. [40] built on video prediction algorithms to predict both future solar irradiance values and sky images. Paletta et al. [24] introduced a video prediction model trained to predict future segmented sky images together with the corresponding solar irradiance values. Logothetis et al. [41] compared the capability of five ASIs to foresee ramp events. Following this, Terrén-Serrano and Martínez-Ramón [42] proposed to extract relevant physical cloud features for solar forecasting from sky images acquired by an infrared imager centred on the Sun. More recently, Nie et al. [43] developed a physics-informed stochastic video prediction model able to generate several possible future sky videos with diverse cloud motion patterns.

Similar techniques have been applied to forecast solar production from satellite data. For instance, Qin et al. [44] utilised a deep learning model to fuse satellite and ground data. Similarly, Bansal et al. [45] proposed a forecasting model based on geostationary observations as well as historic solar and temperature measurements. Si et al. [46] compared time series, NWP and hybrid methods for solar forecasting from satellite observations.

Moreover, sky-image-based models has been shown to outperform algorithms using satellite observations in intra-hour deterministic and probabilistic solar irradiance forecasting [27]. This can be partly explained by the ability of sky images to not only capture rapidly varying information such as incoming clouds visible from the camera, but also contextual information including the current weather conditions, the position of the sun, the types of cloud and aerosol properties. Consequently, although sky cameras are intrinsically unable to provide truly accurate forecasts beyond a certain horizon based on their field of view and the speed of clouds, the contextual spatio-temporal features they can provide have the potential to benefit long-term predictions including satellite-based forecasts.

1.2. Limitations

Despite showing a strong potential to become a prevailing approach in the future, the utilisation of deep learning models for solar forecasting faces several limitations [47] such as a recurrent temporal lag penalising critical event prediction [39], a lack of transparency of current architectures, as well as a need for improved uncertainty quantification methods [48, 49, 27, 50, 43].

Another key challenge of data-driven methods is the difficulty to generalise knowledge acquired at one location to a new place. Even though data augmentation methods have been shown to improve generalisation skills [51, 52], the learning of deep models is bounded by the diversity of the training samples in terms of weather patterns, sky conditions as well as imaging setups. In that respect, preliminary studies applying direct transfer learning to unseen data have shown a significant performance drop due to the distribution mismatch between the datasets [53, 54].

1.3. Transfer learning

Conditioned on the availability of local data, supervised transfer learning methods offer the possibility of a local fine-tuning (i.e. calibration) based on a target data dis-

tribution to improve the generalisation skills of an expert model trained on a different (but related) source dataset that often contains a larger amount of data. In particular, the large quantity of open source sky image datasets is a significant opportunity for the development of such transfer learning techniques [55]. In that respect, several transfer learning approaches based on three datasets collected on different continents have been investigated in Nie et al. [53]. These techniques were compared in a range of data availability contexts from the absence of data (i.e. no local sky observation) to data scarcity (i.e. a few days of atmospheric observations) and data sufficiency (i.e. several years of data) [56].

Although the study by Nie et al. [53] highlights the poor generalisation skills of solar forecasting models when applied to a different location (set-up also known as zero-shot learning : ZSL), transfer learning seems to provide more encouraging results: by learning diverse features over a combination of multiple datasets, models fine-tuned on target data achieve superior performances compared to local baselines. In particular, the outcome of the study showed a significant potential of such methods for model calibration when limited local data are available.

This specific task described as *few-shot learning* (FSL) is especially relevant when the availability of local sky image data is constrained for various reasons, e.g. starting of a new solar site, setup of a novel camera, modified imaging setup, changing solar production capacity, recent changes in the shading configuration or soiling events. In these conditions, the rapid adaptation of the forecasting algorithm is challenging but key for short-term operations.

1.4. Objectives of the study

The main goal of this work is to evaluate and improve the application of already trained deep solar forecasting models to unseen locations such as a new solar farm. To improve this cross-site generalisability, the application of pre-trained neural networks to a new site without fine-tuning (i.e. ZSL) or via a moderate calibration (i.e. FSL) based on as little as one day of data is explored. The proposed method improves over Nie et al. [53] in data scarce conditions by applying a simple but highly effective physics-informed normalisation strategy for solar data based on a clear-sky model. In particular, this study aims at addressing the following questions :

- Can the normalisation of PV and irradiance data via clear-sky models improve transfer learning for solar forecasting?
- Which weather conditions provide the most valuable information for local fine-tuning via FSL?
- To what extent can data augmentation facilitate transfer learning?
- How do models trained on a single day of data using FSL compare with local and fine-tuned models based on a large dataset?

To tackle these problems, the present study compares the proposed FSL/ZSL approaches with the transfer learning baselines introduced in Nie et al. [53]. The chosen learning method consists of adapting a solar power forecasting neural network pre-trained on a large sky image dataset collected at SIRTA atmospheric observatory near Paris [57]

(sky images and global horizontal irradiance measurements) to a target dataset generated at Stanford University [58] (SKIPP'D dataset: sky images and PV output values) with no or limited local target data for fine-tuning.

The rest of the article is structured as follows: Section 2 describes in more details the two datasets utilised in the experiments including sky images, irradiance and PV output measurements. Following this, Section 3 introduces the proposed transfer learning strategies including ZSL and FSL methods. Finally, Sections 5 and 6 summarise and discuss the main results obtained in the study.

2. Data and preprocessing

The study focuses on the operational challenge of providing accurate solar power forecasts to newly-opened solar facilities, which have not yet generated enough data to train a reliable deep learning model [53] but which have access to a large quantity of label data collected at other sites. To address this challenge, transfer learning approaches aim at adapting a machine learning model pre-trained on a (larger) source dataset (i.e. SIRTA) to a (smaller) target dataset (i.e. SKIPP'D).

2.1. Source dataset: SIRTA

The source dataset used to pre-train solar forecasting models is collected at SIRTA, an atmospheric observatory part of the Baseline Surface Radiation Network [59] and located near Paris in France. The dataset is composed of about 450,000 samples that were generated between 2017 and 2019 at a 1- to 2-minute resolution (1 min in 2017, 2 min in 2018-2019).

Irradiance measurements. The site is equipped with diverse irradiance measurement devices including a pyrheliometer (downwelling direct solar irradiance), pyranometers (downwelling diffuse and global solar irradiance, upwelling global solar irradiance) and pyrgeometers (downwelling and upwelling infrared irradiance). The downwelling global horizontal solar irradiance (GHI) used in this work corresponds to the amount of solar energy received by a 1m^2 surface parallel to the ground over a second (unit: W/m^2). The GHI is measured at a 1-second resolution on site by a ventilated pyranometer (Kipp and Zonen CM22) and reported as its per minute average. This sensing device has been operating on site since October 2005. To ensure reliable measurements, the pyranometer is regularly calibrated by the World Radiation Center located in Davos and cleaned from 3 to 5 times a week.

Sky images. Cloud cover observations are collected by an EKO SRF-02 hemispherical sky camera located near the pyranometer. Figure 1A depicts a sky image taken by this all-sky imager (ASI). From the original resolution of 768×1024 pixels, sky images are centred and cropped to remove some of the side uninformative black frame and rescaled to a more standard 64×64 pixel resolution for further analysis by neural networks. This processing step includes a low pass filtering to prevent aliasing effects induced by down-scaling [60].




Figure 1: Sky images taken at (A) SIRTA atmospheric laboratory and (B) Stanford University. Source: Google Map and Cros et al. [61]. The pyranometer and solar panels used to collect solar power data are highlighted in yellow in the figure. In subplot (A), the pyranometer and sky camera are 115 meters away, and in subplot (B), the studied solar panels are located about 125 meters away from the camera.

2.2. Target dataset: *SKIPP'D*

The target dataset used to fine-tune and test solar forecasting models is generated at Stanford University in California, US [62]. The 140,000 samples utilised in this work were collected from March 2017 to November 2019 at a 1 minute resolution.

PV output measurements The PV output data are collected from solar panel arrays approximately 125 meters away from the camera, situated on the top of the Jen-Hsun Huang Engineering Center at Stanford University, with an tilt angle of 22.5° and an azimuth angle of 195° . The PV panels are manufactured with poly-crystalline technology and the system is rated at 30.1 kWp. The PV power generation data are logged with 1-min frequency and are minutely averaged. The forward average is applied, e.g. value at 8:00:00 am representing the average PV generation from 8:00:00 to 8:00:59 am.

Sky images The sky images are frames from videos recorded during daytime (6:00 AM \sim 8:00 PM ¹) by a 6-megapixel 360-degree fish-eye camera (Hikvision DS-2CD6362F-IV2) located on top of the Green Earth Sciences Building (37.427° , -122.174°) at Stanford University and oriented towards 14° south by west. The camera captures video with a resolution of 2048×2048 pixels at 20 frames per second (fps) and the images (*.jpg*) are extracted from the video at 1-min sampling frequency and are down-sampled to a resolution of 64×64 pixels.

2.3. Solar measurement normalisation via clear-sky models for improved inter-site data compatibility

In this study, we normalise GHI and PV output based on their theoretical values under clear sky condition. The resultant ratio is called clear-sky index k_c (CSI), which can be expressed by Equation 1 below:

$$k_c = \frac{I}{I_{\text{clr}}} \quad \text{or} \quad \frac{PV}{PV_{\text{clr}}} \quad (1)$$

¹Data were recorded based on the local time zone, which is either Pacific Standard Time (PST) or Pacific Daylight Time (PDT).

Where the PV and I are the actual PV power production and irradiance measurements, respectively, and P_{clr} and I_{clr} are the corresponding theoretical clear-sky estimates.

The CSI reflects the impact of the cloud cover on incident solar irradiance and, by extension, PV power generation [63]. According to Equation 1, the actual solar level can be easily retrieved based on their clear-sky estimates (details on estimating the clear-sky values of GHI and PV output can be found in the following paragraphs). Therefore, instead of predicting the absolute solar value (I or PV), the CSI can be used as a proxy to focus the modelling on the spatio-temporal dynamics of the cloud cover by factorising the clear-sky solar flux (see Section 3.3). In addition, the CSI provides a normalisation approach that might ameliorate generalisation skills of deep learning models (e.g. application of a model to a different solar site) and facilitate dataset compatibility for multi-location data modelling approaches including data fusion and transfer learning [53].

The estimation of clear-sky GHI and PV output is primarily geometric, hence can be solved based on the location of the observer on Earth. For the SIRTAsite, a clear-sky model external to this work is used (McClellan, [64]). For the Stanford PV facility, the clear-sky solar power output PV_{clr} is modelled based on sun angles and PV panel orientations as shown in the following Equation 2 [65, 66]:

$$\mathcal{P}_{\text{clr}}(t) = I_m A_e \{ \cos \epsilon \cos \chi(t) + \sin \epsilon \sin \chi(t) \cos[\xi(t) - \zeta] \} \quad (2)$$

Where I_m is the maximum solar irradiance, 1000 W/m²; A_e is the effective PV panel area, 24.98 m², which is obtained from a least square fit with the real panel output of 12 clear sky days (details can be found in study by Sun et al. [37]); ϵ and ζ are elevation and azimuth angles of the solar PV arrays, which are 22.5° and 195°, respectively; $\chi(t)$ and $\xi(t)$ are the zenith and azimuth angle of the sun, which can be estimated for any minute of the year from the empirical functions [65].

Additionally, the CSI is used to define the smart persistence model (SPM) based on the expected clear-sky irradiance changes over the forecast horizon T (Equations 3 and 4). The proposed modelling approaches are compared with this baseline model (see Section 4 for more details).

$$I_{\text{SPM}}(t + T) = I_{\text{clr}}(t + T) \frac{I(t)}{I_{\text{clr}}(t)} \quad (3)$$

$$PV_{\text{SPM}}(t + T) = PV_{\text{clr}}(t + T) \frac{PV(t)}{PV_{\text{clr}}(t)} \quad (4)$$

3. Methodology

3.1. Deep learning architecture

The deep learning model used in this study is the ConvLSTM model (Figure 2) presented in [39] in a solar forecasting benchmark and adapted in [53] for the initial transfer learning study. The model is composed for parallel encoders for past solar measurements over 15 minutes at a 2 minute resolution (last eight lag terms of GHI or PV output) and past sky images (last eight RGB images). The encoder for auxiliary data is composed of a set of densely connected layers whose sequential output is processed by a recurrent neural network (long short-term memory network: LSTM). Similarly, input

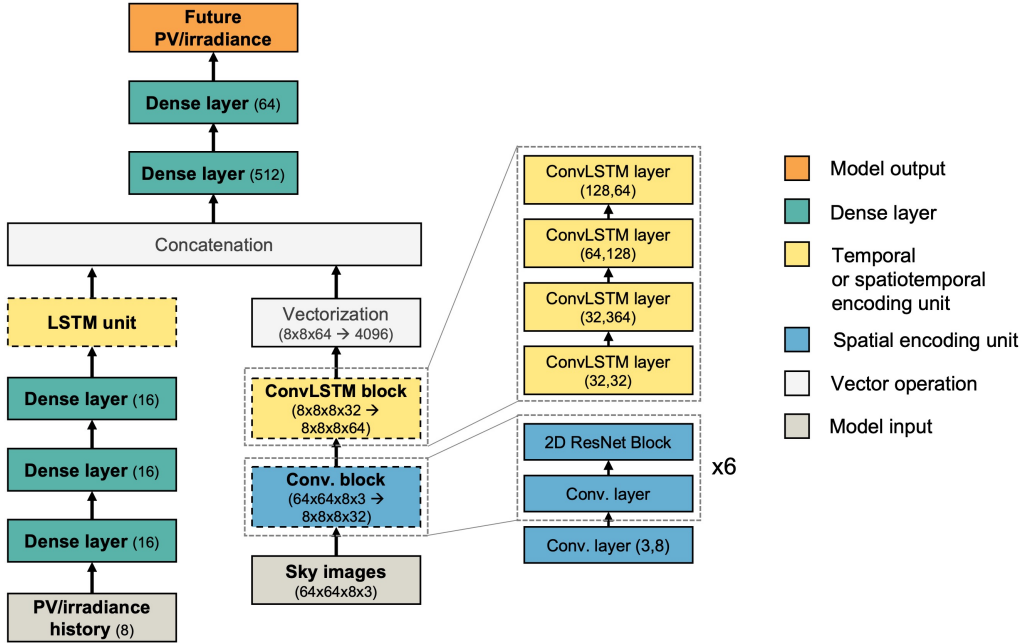


Figure 2: Deep learning architecture used to forecast solar energy. Adapted from [39] and [53].

sky images are first individually encoded by 2D convolutional layers (spatial encoding) followed by a 2D recurrent neural network (convolutional LSTM). The output of both encoders is merged (concatenation) prior to be fed to another set of dense layers with decreasing size: 512, 64 and finally 1 neuron corresponding to the solar prediction. The lead time considered in this study is 15 minutes. Overall, the model is composed of 4.36M trainable parameters.

3.2. Transfer learning approaches

This section presents a range of transfer learning approaches for applying a pre-trained solar forecasting neural network to a different site following on the diverse data availability scenarios.

Zero-shot learning (i.e. no local sky observations) The primary objective of this task is to evaluate the potential of deep solar forecasting models to generalise beyond the location they have been trained on. The experimental approach consists of training a model on a source dataset (SIRTA) and to directly test it on another dataset (SKIPP'D) without local fine-tuning (ZSL). This zero-shot regression task was initially addressed in [55] by normalising both solar time series based on robust statistics (e.g. standard deviation, 95% quantile). As such, the model was trained to forecast the normalised solar level that was scaled back to the appropriate range (e.g. about 0 to 30kW for PV output or 0 to 1300W/m² for irradiance measurements). This method resulted in significantly worse local performances than the ones of models locally trained. The authors observed




Figure 3: Illustration of the FSL approach. Adapted from [53].

that the majority of the corresponding forecast error originates from a poor scaling for the output variable. Indeed, the statistics used in this method partly depend on location-specific parameters (e.g. the sun position, weather conditions) and on the characteristics of the measuring device (pyranometers *versus* power meters).

To address that limitation, a proposed method consists of predicting the CSI, i.e. the solar value normalised by the corresponding clear-sky estimation. This method is expected to be a more robust approach than using local data statistics, especially in the context of limited data availability.

Few-shot learning (e.g. one day of data). This field of machine learning focuses on solving tasks based on limited data. This often requires a large and diverse auxiliary dataset to learn relevant features that can be reused for the target task via fine-tuning (Figure 3).

In this study, ‘limited data’ corresponds to one day of measurements. In practice, a range of neighbouring days are selected from the target dataset (SKIPP’D, US) for fine-tuning a source model trained on about half a million samples of cloud cover observations collected over three years in another location (SIRTA, France). These training days correspond to diverse sky types used to evaluate the benefit of different meteorological conditions for model calibration: clear-sky, scattered clouds, broken-sky, overcast to clear-sky, overcast (Figure 4). To describe a realistic scenario, the associated validation set is composed of 4 days (three clear-sky days and one cloudy day) occurring closely after the selected training days (Figure 5).

Fine-tuned model on the extended local dataset (269 days). The ZSL and FSL approaches are compared with the transfer learning configuration based on the full target dataset for calibration [53]: the model pre-trained on SIRTA dataset is fine-tuned on 269 days of sky images collected at Stanford University. The performance of the resulting algorithm will provide some insights on the progress margin of models based on no or limited data.




Figure 4: Days with diverse weather conditions used for local fine-tuning via FSL: clear-sky day (A), scattered clouds (B), broken cloud cover (C), overcast to clear-sky (mix) (D), overcast (E). Time windows containing night times and missing data have been truncated.




Figure 5: Days used for validating the models fine-tuned via FSL. Discontinuities in the time series correspond to missing data. Time windows containing night times and missing data have been truncated.

3.3. Modelled output variable

In the context of transfer learning, multitask learning, and multi-site dataset integration, the choice of the output variable being predicted by the model can have a strong impact on the overall forecasting performance. This study investigates three methods to address this strong limitation (Figure 6).

Time series normalisation using local data statistics. To deal with the heterogeneity of the modelled variables from the source and target datasets (e.g. solar irradiance and PV output) in terms of scale and distribution, a first approach introduced by Nie et al. [53] consists of normalising each time series independently by the same local statistics to facilitate dataset matching. In practice, the model is trained on normalised time series (Equation 5), which can be scaled back to the original space by reversing the normalisation function. In this work, the GHI and PV output time series are normalised by the respective 95% percentile, 853.0 W/m² and 23.8 kW respectively (Figure 7). It should be noted that this method — Time series normalisation using local data statistics (TSN) — requires reliable data statistics to compute the normalised solar value (NSV) and that it is prone to biases induced by the distribution of location-specific weather patterns (e.g. clear-sky *versus* overcast) and the sampling distribution over the collection period (e.g. summer *versus* winter). This limiting factor is strengthened in the context of local data scarcity. As such, additional experiments will highlight the impact of the local normalising constant (i.e. 95% percentile of the PV output time series) on the overall forecasting skill. The training loss for TSN method is defined by Equation 5:

$$L_{NSV} = \frac{1}{N} \sum_{k=1}^N (\widehat{NSV}_{t+15\text{min}} - NSV_{t+15\text{min}})^2 \quad (5)$$




Figure 6: Approaches addressing the output variable heterogeneity (e.g. solar irradiance *versus* PV output, different locations or sampling periods): (1) time series normalisation using local data statistics based on the normalised solar value (NSV), (2) normalisation via the clear-sky index (CSI), and (3) clear-sky index errors weighing based on the clear-sky solar value (CSV).

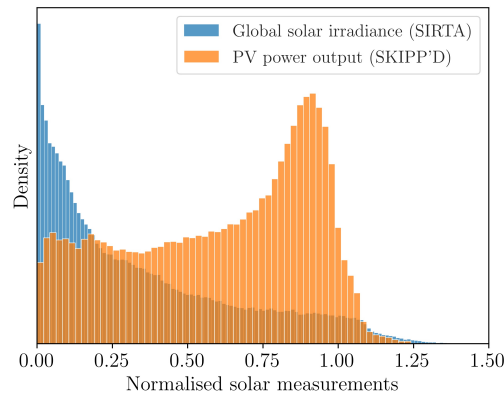


Figure 7: Distribution of solar irradiance and PV output measurements normalised by the 95% quantile for distribution matching.




Figure 8: Distribution of the CSI (k_c) and absolute solar values for SIRTA (GHI) and Stanford (PV output) training sets.

Normalisation by the clear-sky solar value. Another approach introduced in the literature consists of modelling the CSI based on a local clear-sky model instead of the absolute solar value [37, 67, 68]. In addition to past sky images, the model is fed past CSI values instead of past GHI (or PV) measurements and is trained to predict the future CSI. As such the resulting loss function is based on CSI values (Equation 6). The future clear-sky irradiance or clear-sky PV output is then used to convert the CSI into the variable of interest (i.e. GHI or PV output). Contrary to the TSN method, this type of normalisation approach does not require local data statistics. Additionally, the model does not have to handle the scale of the target distribution as it is modelled by the clear-sky model. This limitation was shown to cause a significant performance loss in low data availability conditions [53, 54].

$$L_{\text{CSI}} = \frac{1}{N} \sum_{k=1}^N (\widehat{k_{c,t+15\text{min}}} - k_{c,t+15\text{min}})^2 \quad (6)$$

Weighted clear-sky index errors. A possible limitation of the CSI method is the changing target distribution induced by the normalisation (Figure 8). As such, the model will uniformly focus on samples regardless of the magnitude of the corresponding GHI or PV output value (Figure 11 in [37]). As a result, the model’s focus might partly shift from high clear-sky solar conditions corresponding to high absolute errors towards low clear-sky solar conditions (i.e. beginning and end of the day). The resulting absolute solar forecasting performance might therefore drop.

To tackle this challenge while benefiting from the CSI, an alternative consist of weighing the output of the model by the corresponding clear-sky solar value (CSV) such that the objective function L is based on the absolute solar value (ASV) instead of the CSI as $ASV = k_c \times CSV$. The resulting loss function based on the weighted clear-sky index

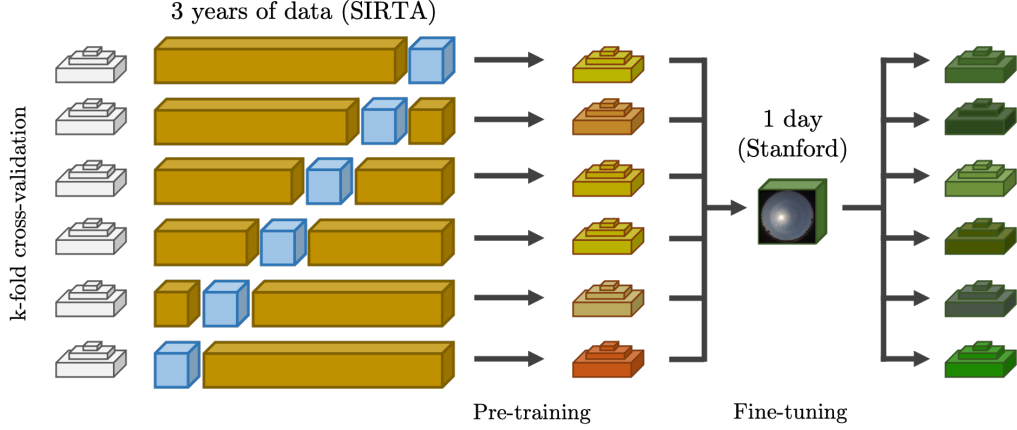


Figure 9: FSL pipeline: The sub-models are first trained via k -fold cross-validation (brown: training set, blue: validation set) on the source dataset (SIRTA) prior to calibration on a single day of data collected at the target location (Stanford).

errors (WCE) depicted in Equation 8 weighs high top-of-atmosphere solar conditions more while the algorithm models the CSI.

$$L_{ASV} = \frac{1}{N} \sum_{k=1}^N ([\widehat{k_{c,t+15min}} - k_{c,t+15min}] \times CSV_{t+15min})^2 \quad (7)$$

$$= \frac{1}{N} \sum_{k=1}^N ([\widehat{k_{c,t+15min}} \times CSV_{t+15min} - ASV_{t+15min}])^2 \quad (8)$$

4. Results

The proposed experiments are implemented using the deep learning framework `Pytorch`. When the training dataset contains enough data for pre-training and fine-tuning (i.e. full SIRTA and SKIPP'D dataset), ten-fold cross-validation is applied during the learning phase. The predictions of the resulting 10 sub-models are averaged during the model evaluation phase. In the context of data scarcity (i.e. less than 10 days), the 10 sub-models trained on the source dataset are all fine-tuned on the same small training set and evaluated on a validation set composed of distinct days (Figure 5). For instance, in the context of FSL, the 10 sub-models are first trained via ten-fold cross-validation on SIRTA dataset. Then, each sub-model is fine-tuned on a single day (e.g. clear-sky day) and validated on the four days composing the validation set (Figure 9). All models trained in study are tested on the same test set composed of 10 clear-sky days (5359 samples) and 10 cloudy days (4292 samples). No statistical post-processing is performed on the output of the model.

Evaluation metrics comparing model predictions $\widehat{y_{t+15min}}$ with ground truth $y_{t+15min}$ include the root mean square error (RMSE - Equation 9) and the mean bias error (MBE -

Equation 10). In addition, the forecast skill (FS) metric is used to quantify the improvement of the models relative to the smart persistence model (SPM) based on the RMSE (see Equation 11). Furthermore, a distribution-oriented approach is used for forecast verification via scatter plots [69].

$$\text{RMSE} = \frac{1}{N} \sum_{k=1}^N (\widehat{y_{t+15\text{min}}} - y_{t+15\text{min}})^2 \quad (9)$$

$$\text{MBE} = \frac{1}{N} \sum_{k=1}^N (\widehat{y_{t+15\text{min}}} - y_{t+15\text{min}}) \quad (10)$$

$$\text{Forecast Skill} = \left(1 - \frac{\text{RMSE}_{\text{model}}}{\text{RMSE}_{\text{SPM}}}\right) \times 100 \quad (11)$$

In addition, the results obtained in low data availability conditions are compared with traditional training approaches based on a large quantity of data (e.g. 2 to 3 years). The resulting performance of both local learning (i.e. training on the full Stanford dataset) and transfer learning (i.e. pre-training on SIRTA dataset and fine-tuning on the full Stanford dataset) strategies are depicted in Table 1. Furthermore, data augmentation being a key solution for FSL to increase the diversity of a small training set, the impact on the forecasting skill of random rotations ($0 - 360^\circ$) applied to sky images during training is also presented in Figure 10. This data augmentation technique was shown to improve the accuracy of solar irradiance predictions in previous studies but it has never been tried with PV output forecasting [51]. It is specifically utilised here to highlight the general potential of augmenting the data in data scarce conditions. In practice, there exist a wide range of other data augmentation strategies that have been applied to sky images, specifically vertical, horizontal and temporal flips, Gaussian noise, Colour casting, brightness adjustment, SMOTE, Mixup [51, 52]. Combining all these methods and others during training would likely ameliorate the solar forecasts further while improving the generalisation skills of the models to specific contexts such as lens soiling, cameras in different hemispheres, and different image capturing settings.

4.1. Expert models

The first results presented in Table 1 highlight the forecasting accuracy of the deep learning model following three training approaches: local training (SKIPP'D), local fine-tuning of a pre-trained model based on SIRTA data, and local fine-tuning of a pre-trained model based on SIRTA data using data augmentation. In addition, the three modelling strategies described in Section 3.3 are tested: TSN, CSI and WCE. This preliminary experiment aims to set competitive baselines — expert models — using the entire source and target datasets (SKIPP'D and SIRTA). These algorithms will be compared with ZSL and FSL methods in the following sections.

In terms of average forecasting error, the models based on the clear-sky index (CSI or WCE) often outperforms the TSN method in cloud-free conditions. For instance, the forecasting skill of the local model trained via TSN without data augmentation is around 31.0% in clear-sky days compared to 53.7% and 49.4% for the CSI and WCE methods. In cloudy conditions, the various methods perform similarly with a forecast skill metric




Figure 10: Original sky image and corresponding clock-wise rotations by 90, 180 and 270 degrees.

Table 1: Performance of ‘expert’ solar power forecasting models (15-min lead time) based on a large target training set (SIRTA) and three training strategies: time series normalisation (**TSN**), predicting the clear-sky index (**CSI**), and weighing the clear-sky index errors (**WCE**) with the corresponding CSV to generate the ASV from the CSI prediction ($ASV = k_c \times CSV$). The forecasting performance of these expert models are provided for comparison with ZSL (Table 2) and FSL (Table 3) methods.

Training strategies	RMSE ↓ [W/m ²] (Forecast Skill ↑ [%])		MBE	
	Cloud-free	Cloudy	Cloud-free	Cloudy
SPM	0.36 (0%)	4.79 (0%)	0.11	-0.17
Local training				
- TSN [53]	0.25 (31.0%)	3.87 (19.1%)	0.12	0.33
- TSN (+ Rotations)	0.18 (50.6%)	3.92 (18.0%)	0.11	-0.17
- CSI	0.17 (53.7%)	3.91 (18.3%)	0.04	0.32
- CSI (+ Rotations)	0.15 (59.5%)	3.84 (19.8%)	0.02	0.48
- WCE	0.18 (49.4%)	3.90 (18.4%)	0.04	0.27
- WCE (+ Rotations)	0.16 (54.7%)	3.86 (19.4%)	0.02	0.31
pre-training on SIRTA data				
- TSN [53]	0.21 (40.7%)	3.92 (18.1%)	0.10	0.39
- TSN (+ Rotations)	0.17 (52.7%)	3.87 (19.1%)	0.02	0.36
- CSI	0.18 (51.2%)	3.91 (18.3%)	0.07	0.58
- CSI (+ Rotations)	0.16 (56.9%)	3.85 (19.5%)	0.02	0.50
- WCE	0.23 (34.7%)	3.90 (18.5%)	0.07	0.41
- WCE (+ Rotations)	0.15 (57.4%)	3.87 (19.1%)	0.02	0.33
pre-training on SIRTA data with data augmentation				
- TSN [53]	0.22 (39.8%)	3.98 (16.8%)	0.10	0.23
- TSN (+ Rotations)	0.19 (47.7%)	4.04 (15.6%)	-0.02	0.17
- CSI	0.23 (37.0%)	3.91 (18.3%)	0.14	0.29
- CSI (+ Rotations)	0.15 (57.7%)	3.87 (19.1%)	0.04	0.53
- WCE	0.18 (50.4%)	3.93 (17.8%)	0.04	0.37
- WCE (+ Rotations)	0.16 (55.7%)	3.89 (18.8%)	-0.01	0.42

Table 2: Test set performance of the ZSL approach for solar power forecasting models (15-min lead time) based on a large source training set (no local fine-tuning) and three training strategies: time series normalisation (**TSN**), predicting the clear-sky index (**CSI**), and weighing the clear-sky index errors (**WCE**) with the CSV to generate the ASV from the CSI ($ASV = k_c \times CSV$). As shown by the MBE in clear-sky conditions, the difficulty to estimate the solar power output from the current sky image results in negative forecast skill scores (i.e. the models under-perform the SPM). In most cases, however, the RMSE obtained in cloudy conditions outweighs the one corresponding to cloud-free days.

Training strategies	RMSE ↓ [W/m ²] (Forecast Skill ↑ [%])		MBE	
	Cloud-free	Cloudy	Cloud-free	Cloudy
SPM	0.36 (0%)	4.79 (0%)	0.11	-0.17
- TSN [53]	4.07 (-1033.5%)	5.41 (-13.1%)	-3.19	-3.17
- TSN (+ Rotations)	2.81 (-683.7%)	5.83 (-21.8%)	-0.92	-3.37
- CSI	2.17 (-504.7%)	4.28 (10.5%)	-0.82	-0.68
- CSI (+ Rotations)	1.68 (-367.0%)	4.27 (10.69%)	0.52	-0.83
- WCE	2.20 (-512.6%)	4.51 (5.7%)	-0.35	-1.20
- WCE (+ Rotations)	1.55 (-332.7%)	4.24 (11.4%)	0.74	-0.53

ranging from 15.6 to 19.8%. Applying random rotations to input sky images to augment the training dataset leads to consistent gain for cloud-free samples and for methods based on the CSI in cloudy conditions. This shows that this type of data augmentation can also be beneficial for forecasting the power output of a solar site even though the orientation of tilted solar panels — as opposed to flat sensors such as a pyranometers — hinders the rotational invariance of the problem. We hypothesise that the model might be able to recognise the orientation of the rotated image based on static objects in the image (e.g. trees or buildings) or that the benefit of showing more diverse spatiotemporal patterns of the cloud cover outweighs the difficulty to learn the specific solar panels orientation. Note that the forecast error associated with cloud-free days is an order of magnitude lower than the one corresponding to cloudy days. The mean square error loss function amplifies this difference during training, hence it predominantly stirs the model attention towards samples with clouds, which are more challenging. As a results, the impact of cloud-free samples is often marginal in an overall training objective based on diverse cloud cover conditions.

4.2. Zero-shot learning results

The results of the ZSL experiments are shown in Table 2. To begin with, the performances of the models in cloud-free conditions are significantly worse than the SPM with forecast skills ranging from -1033.5% to -332.7% . In these conditions, the error associated with clear-sky days is comparable to the one obtained in cloudy days. Note that the magnitude of the corresponding MBE is significantly higher than in Table 1. For instance, the TSN gives a MBE of about $-3.2W/m^2$ for a RMSE of $4.07W/m^2$ in cloud-free conditions and $5.41W/m^2$ in cloudy conditions. For most configurations, the use of data augmentation via rotations provides a performance gain, especially when the sky is clear. The use of the CSI in the modelling (CSI and WCE methods) gives a significant

Table 3: Performance of the FSL approach for solar power forecasting models (15-min lead time) based on a large source training set and a small local dataset (1 day) via three training strategies. Similarly to the ZSL results, the SPM based on the last solar measurement outperforms FSL methods in cloud-free conditions (i.e. negative forecast skills). In cloudy conditions, however, positive skills are observed together with error values outweighing the ones obtained in clear-sky conditions based on a quadratic loss.

Training strategies	RMSE ↓ [W/m ²] (Forecast Skill ↑ [%])		MBE	
	Cloud-free	Cloudy	Cloud-free	Cloudy
SPM	0.36 (0%)	4.79 (0%)	0.11	-0.17
Time series normalisation [53]				
- Clear-sky	1.16 (-222.9%)	4.20 (12.2%)	0.22	0.87
- Scattered clouds	1.59 (-341.8%)	4.55 (5.0%)	0.24	1.82
- Broken-sky	2.01 (-459.2%)	4.21 (12.1%)	0.80	0.91
- Overcast to clear-sky	1.36 (-278.5%)	4.14 (13.4%)	0.25	0.91
- Overcast	2.04 (-468.0%)	4.24 (11.4%)	-0.32	-0.56
Clear-sky index				
- Clear-sky	2.18 (-506.2%)	6.04 (-26.2%)	1.82	4.32
- Scattered clouds	2.03 (-466.6%)	6.88 (-43.8%)	1.58	5.16
- Broken-sky	1.99 (-452.7%)	4.54 (5.17%)	1.57	1.97
- Overcast to clear-sky	2.37 (-560.9%)	5.11 (-6.8%)	2.07	3.01
- Overcast	1.42 (-296.1%)	4.03 (15.75%)	0.78	0.42
Weighted clear-sky index errors				
- Clear-sky	2.17 (-504.3%)	6.2 (-29.6%)	1.82	4.49
- Scattered clouds	2.25 (-527.7%)	6.16 (-28.7%)	1.80	4.27
- Broken-sky	2.16 (-501.7%)	4.91 (-2.7%)	1.74	2.52
- Overcast to clear-sky	2.26 (-530.3%)	5.52 (-15.3%)	1.91	3.32
- Overcast	1.53 (-325.9%)	4.20 (12.1%)	0.69	0.30

benefit compared to the TSN approach. For example, the forecast skill associated with the WCE strategy reaches 11.4% in cloudy conditions compared to around 17 – 19% for the models trained on the entire target dataset (Table 1). This performance gap with the TSN method can be explained by the difficulty with this modelling strategy to adjust the scale of the output variable [53] resulting in a large MBE (Table 2). On the contrary, modelling approaches parameterised by the CSI (e.g. CSI, WCE) have the scale right by design: clear-sky conditions correspond to a model prediction of about 1 regardless of the actual output variable distribution and scale. This analysis advocates for the use of the CSI in a ZSL context, especially when the variables of interest have widely different scale (e.g. solar irradiance *versus* solar power output, solar irradiance at different latitudes).

Overall, the WCE strategy with data augmentation offers the best ZSL performance in terms of forecasting skill. However, improving the zero-shot-regression performance in clear-sky conditions remains crucial to close the gap with expert models learning from a large amount of target samples, and thus providing more reliable local predictions.

4.3. Few-shot learning results

Table 3 describes the accuracy of the models in the FSL setting, i.e. data scarce conditions. The pre-trained source models presented in Section 4.1 are fine-tuned based on a single day of cloud observations using data augmentation. Interestingly the modelling strategy (e.g. TSN, CSI and WCE) give widely different results based on the type of weather used for model fine-tuning. For instance, clear-sky images (clear-sky and overcast to clear-sky days) are the most informative when the predicted output variable is the solar measurement normalised by the 95% quantile. This result can be explained by the conclusion of Section 4.2, i.e. the model’s difficulty to predict the scale of the output variable via TSN (i.e. systematic bias) is partly overcome by observing the sun in sky images (i.e. clear-sky and overcast to clear-sky samples). As a result, the model calibrates its predictions by refining the learnt sun-position-to-power-output equation based on local observations of the sun and corresponding power output measurements. Similarly, observing overcast samples is not as efficient to learn the magnitude of the output variable in clear-sky conditions, which leads to poorer performances. Note that the application of random rotations to training clear-sky image results in a significantly wider coverage of sun positions seen by the model compared to a training without data augmentation, which only describe a single sun trajectory over a day.

In contrast, models predicting the CSI can easily learn the scale of the clear-sky PV output distribution, i.e. from 0 to about 1.3. However, the magnitude of solar power production when the sun is occluded remains challenging as it depends on the type of sensor (e.g. solar panel, pyranometer) and its spatial distribution (e.g. solar cell *versus* solar farm). Consequently, learning the left-hand side of the distribution corresponding to samples with an occluded sun, is more informative for the neural networks trained via the CSI. As expected, the best performance is obtained when models are fine-tuned on the overcast day.

Although the cloud-free forecast skill of models trained via FSL is still substantially worse than the one of expert models (i.e. -222.9% *versus* 59.5%), the performance gap narrows in cloudy conditions (e.i., 15.75% *versus* 19.8%). The error associated with cloudy samples outweighing the one of clear-sky samples in practice, the overall performance of the best FSL models is close to the one of expert models. Therefore, we infer that an algorithm modelling the output variable via a clear-sky model (e.g. CSI or WCE strategies) followed by a local calibration on a single day with specific weather conditions (e.g. overcast conditions), offers competitive generalisation skills compared to local models trained on several years of data. Alternatively, the TSN strategy appears to offer more robust results with forecast skill scores ranging from 5 to 13.4% in cloudy conditions, whereas 7 out of 10 forecast skills are negative for the CSI or WCE strategies.

4.4. Overall quantitative comparison

Figure 11 summarises the forecast accuracy of the diverse modelling strategies in different data availability conditions. All the models previously described in this study are reported here including local expert models (i.e. local models and fine-tuned models based on a large target dataset of 522 days), FSL models calibrated via a single day, and ZSL models applied without local fine-tuning. Interestingly, all models based on a large amount of data (‘Local training’ and ‘Transfer learning’ in the figure) perform similarly on average: from 0.15 to 0.25 W/m² in clear-sky days and from 3.85 to 4.05 W/m² in




Figure 11: Comparison of training and modelling approaches based on the forecasting performance of the models on 10 clear-sky days and 10 cloudy days. All results presented in previous sections are illustrated here.

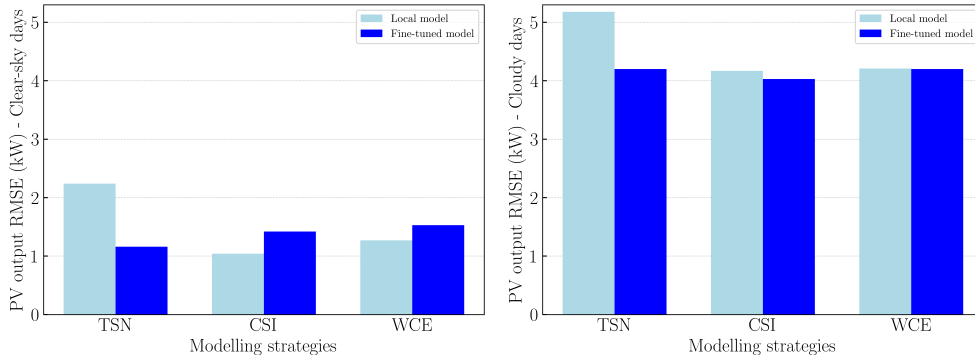


Figure 12: Benefits of model pre-training on a large source dataset (SIRTA) in the context of FSL (i.e. limited availability of target data). For each modelling strategy, the type of weather condition that was shown to provide the best performance in Section 4.3 is used for the local training (i.e. TSN: clear-sky day, CSI & WCE: overcast day). The largest gain is observed for the TSN method in both clear-sky and cloudy days.

cloudy days. In comparison, the error associated with ZSL and FSL approaches is more spread out as illustrated by performance gaps up to 3 W/m^2 among models. However, the difference between expert local models and the best ZSL/FSL methods is narrow, especially in cloudy conditions.

Regarding the gain of model pre-training for FSL, the results widely depend on the type of modelling strategy. Figure 12 highlights the RMSE of the diverse FSL methods based on the weather conditions that were shown to give the most accurate forecasts (Table 3): TSN: clear-sky, CSI & WCE: overcast. Even though the benefits of model pre-training are contrasted for models based on a clear-sky model (i.e. CSI and WCE), a significant improvement is observed for the model trained via TSN in both clear-sky and cloudy days as shown in Nie et al. [53].

4.5. Qualitative observations

Figures 13 and 14 illustrate the forecasting behaviours induced by the various training contexts: local expert model, transfer learning, FSL, and ZSL. Only the best performing model of each category was selected for each training context. Note that the sharp discontinuities particularly visible in Figure 13 correspond to missing data.

For cloud-free days, all algorithms seem to perform similarly (Figure 13). As expected, the ZSL shows a much higher systematic bias. In contrast, the bias induced by the FSL approach is more visible at the beginning and at end of each day. This is likely caused by the poor performance of the clear-sky model in these conditions. In similar conditions, expert models provide more reliable predictions, which indicates that this type of bias was corrected during training. In cloudy weather conditions, the predictions of all models appear highly correlated (Figure 14). This highlights the reliability of ZSL and FSL approaches, and thus the potential of these computer vision algorithms to generalise well to new sites via limited or even no local fine-tuning.

The scatter plots presented in Figure 15 highlight the joint and marginal distributions of the ground truth solar panel output and corresponding model predictions based on

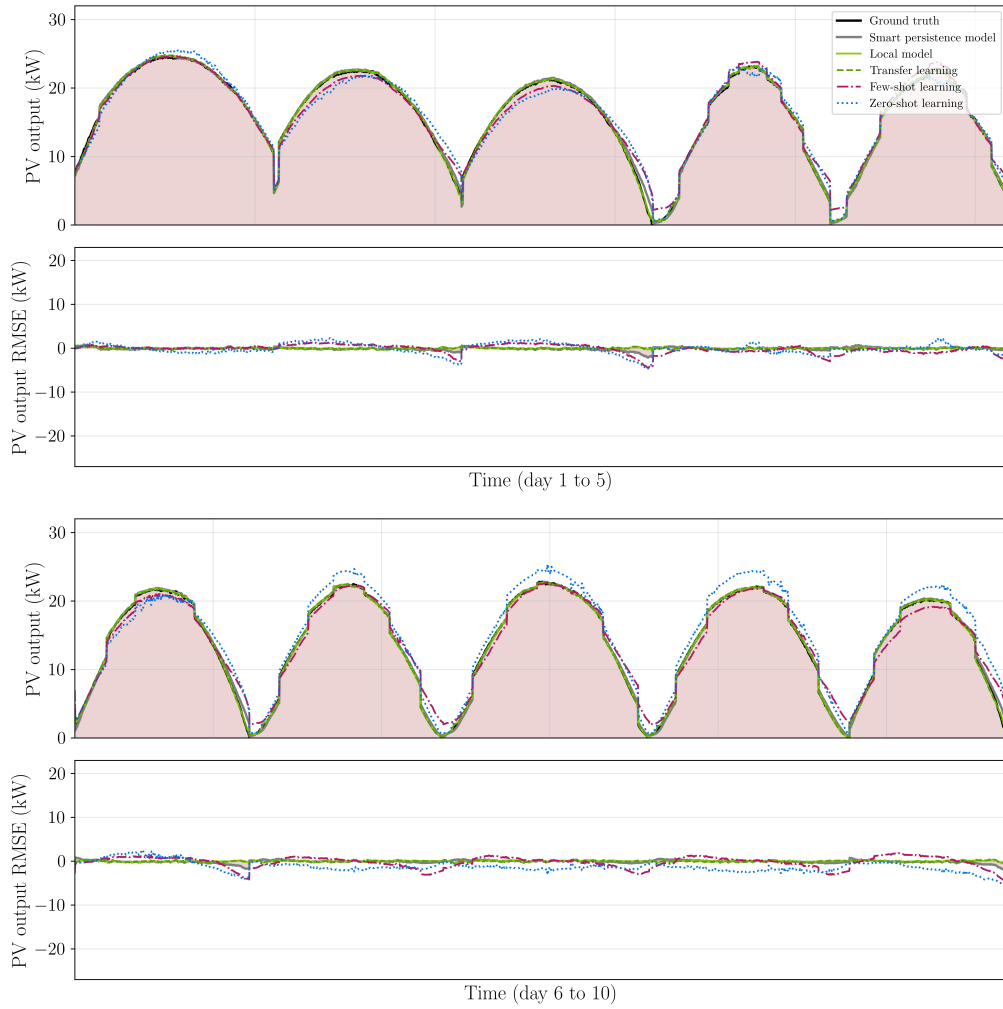


Figure 13: Solar prediction curves of the best target modelling approach for diverse data availability contexts — from no target local data to a large quantity of target samples — based on the forecasting performance of the model on 10 clear-sky days. Discontinuities in the time series correspond to missing data. Time windows containing night times and missing data have been truncated.




Figure 14: Solar prediction curves of the best target modelling approach for diverse data availability contexts — from no target local data to a large quantity of target samples — based on the forecasting performance of the model on 10 cloudy days. Discontinuities in the time series correspond to missing data. Time windows containing night times and missing data have been truncated.




Figure 15: Joint distributions of the best models' predictions in clear-sky conditions (left column) and cloudy weather (right column). From top to bottom : 1. SPM, 2. Zero-shot learning, 3. Few-shot learning, 4. Local model trained on the full dataset via transfer learning. The plots summarise the joint distributions, or equivalently, the marginal and conditional distributions [69].

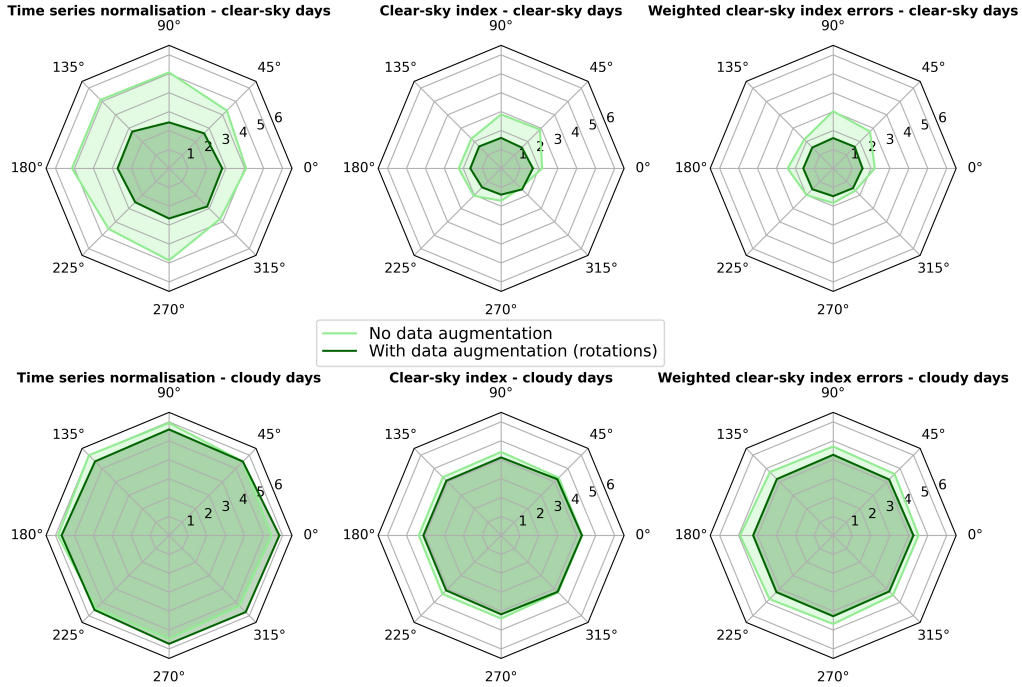


Figure 16: Impact of the target camera (Stanford) orientation on the zero-shot solar forecasting performance quantified by the RMSE for three training strategies. 0° corresponds to the same North-South orientation of the source (SIRTA) and target cameras at test time.

various approaches: SPM, ZSL, FSL, and local experts trained via transfer learning. For clear-sky days, all forecasts appear well correlated with the PV output measurement. As expected, the more local data are used for fine-tuning the better. In cloudy conditions, the distribution of the model predictions is narrower than the one of corresponding observed solar measurements as observed in [39]. Similarly to the first row, the more local training data the better the correlation, even with as little as one day of observations.

4.6. Impact of the camera orientation at the target location

This section evaluates in more details the benefits of augmenting the training data with rotations. Figure 16 describes how the error of the model (RMSE) varies with the orientation of the target sky camera in the context of ZSL. The aim of the experiment is to evaluate the ability of models to transfer to cameras oriented differently or located in a different hemisphere. The results show that the models trained with data augmentation (i.e. rotations from 0 to 360°) are unaffected by the orientation of the target camera. However, the performance of the model trained with standard images only (i.e. no rotations of the source camera), decreases significantly with the rotation of the target camera.

By showing more diverse sky patterns and in particular more varied sun positions in the image, data augmentation via rotations facilitates the generalisation of computer vision models to different camera orientations. This conclusion advocates for the use




Figure 17: Impact of the target normalisation factor on the forecast error associated with the TSN approach (test set). In this particular case, increasing the value of this parameter by about 10% would have decreased the overall RMSE by close to 5%.

of this data augmentation technique when combining several datasets obtained from cameras with diverse orientations.

For the zero-shot regression task, the performances of the CSI and WCE approaches strongly correlate with the accuracy of the underlying clear-sky model. The clear-sky algorithm used for the Stanford dataset was fitted using 12 sunny days [37]. The resulting RMSE is 1.94 kW for 10 test days. This inaccuracy partly explains the difficulty of ZSL methods based on the CSI to reach a RMSE below 1.5 kW in clear-sky days [37]. Results show that with more local training data, the models are able to correct some of the clear-sky model biases (Table 1).

4.7. Normalisation factor

A key component of the TSN approach is the normalisation factor used to match solar data from different sources. In this project the parameter was set to the 95% quantile of the ordered solar measurements — solar irradiance for the SIRTa dataset and PV power output for the Stanford dataset — of the training / validation time series. Although using a quantile improves the robustness of the technique, its value can be biased by various factors: data collection period, type of variable (e.g. PV output, GHI, direct normal irradiance, etc.), local climate, and weather conditions. As a result, an additional tuning of this parameter might be beneficial, especially for the target dataset. Figure 17 illustrates how a change in the value of the target normalisation factor — 23.8 kW in this study — impacts the accuracy of the ZSL method, which is the most affected by this parameter (no model calibration via fine-tuning). For this specific study, we observed that increasing the value of the normalisation factor by 10% could have improved the RMSE by about 5%. Note that the methods based on the CSI are not affected by this limitation.

5. Discussion

This work investigates the ability of computer-vision solar forecasting models to generalise beyond their training location(s). More specifically, two approaches are explored,

namely zero-shot regression learning (i.e. application of a pre-trained algorithm to a new location without local fine-tuning at the target location) and few-shot regression learning (i.e. calibration of a pre-trained model based on limited local data at the target location). Furthermore, the benefit of using a clear-sky model to model the solar variable of interest (solar irradiance or PV power output) is evaluated.

Results show that with the right modelling approach and specific cloud cover data for model calibration, ZSL and FSL methods generalise relatively well with forecast skill scores close to the ones of expert models in cloudy conditions (ZSL: 11.4%, FSL: 15.75%, Local expert: 19.8%). However, the accuracy of these models in clear-sky conditions is worse than the one of the SPM.

For the TSN method, this limitation can be explained by the difficulty to properly match the source and the target distributions: different sensing periods, climates, spatial extents of the sensor (e.g. pyranometer *versus* solar power plant), etc. This matching can be significantly improved via model calibration (via FSL instead of ZSL). Alternatively, the prediction of relative values — CSI/WCE methods or by predicting solar power changes relative to the last measurement [24] — instead of the absolute value would be beneficial. Forecasting a future solar value could indeed be decomposed into two parts: estimating the current level of solar power and predicting the future change relative to the current value. The former challenge is relatively well addressed by neural networks when enough training data is available [70]. However, in scarce data availability conditions, the majority of the forecast error seems to originate from the radiative transfer part. To address this limitation, systematic errors could be corrected locally. Note that the lack of representativeness of a dataset composed of a few subsequent days (FSL) induces some biases that hinder the generalisability of resulting models to other climate patterns such as the ones described in the test sets.

In that respect, the refining of deep learning models via FSL highly depends on the modelling strategy (e.g. TSN, CSI or WCE) and on the weather type of local training data (e.g. clear-sky, overcast, partly-cloudy conditions). This raises the question of the relative value of cloud cover training samples for a given learning objective. With the rapidly growing availability of sky image datasets corresponding to centuries of cloud cover observations [55], training models on such a large quantity of data will be challenging. The development of foundation models trained on diverse tasks and datasets prior to a local fine-tuning is therefore a promising research path. In that respect, transformers [71, 43] and diffusion models [72] are strong candidates due to their capacity to handle large datasets. Furthermore, estimating the potential benefit of each individual training sample could help decrease the size of the training set without degrading the performance. Alternatively, transfer learning techniques aiming at reducing the marginal and conditional feature space mismatch between source and target datasets could provide additional gains tailored to individual data availability conditions.

Regarding the impact of weather on the accuracy of the solar forecasts, cloudy conditions are significantly more challenging than clear-sky days as illustrated by the RMSE and MBE metrics in Table 1. For instance, the RMSE of the SPM is more than 13 times lower during cloud-free days than it is in cloudy conditions (0.36 W/m^2 *versus* 4.79 W/m^2). Furthermore, the forecast skill metric shows that it is more challenging for expert models to improve over the SPM in cloudy conditions than in clear-sky conditions (Table 1). Interestingly, this trend is reversed for few-shot and zero-shot learning methods (Tables 2 and 3) as indicated by a forecast skill ranging from -1033 to -222.9%

in cloud-free conditions and from -43.8 to 15.75% in cloudy conditions. To extend this analysis further, future research could test the performance of the model on more diverse weather types including different categories of clouds and aerosols.

The large discrepancy in terms of RMSE between clear-sky and cloudy days as well as the different solar flux distributions between Paris and San Francisco (Figure 8) highlight not only the dependency of the average prediction accuracy on the weather and cloud type distribution, but also on the local climate conditions. Therefore, the future development of multi-location datasets [73] will allow to characterise the impact of diverse climate conditions on the solar forecasting performance of deep learning models.

Augmenting the local training dataset with sky image rotations is another valuable strategy even when applied to solar power output forecasting, which is typically not fully invariant by rotation. Interestingly, the gain offered by data augmentation is significantly higher in data scarce conditions compared to local expert model conditions (Tables 1 and 2). This indicates that augmenting the data further would likely close the performance gap between both types of configurations. Beside other data augmentation techniques, simulating local data for model calibration could provide further forecasting gains. This could be achieved by generative models such as generative adversarial networks or diffusion models [21].

6. Conclusion

Accurate solar forecasting is crucial to address the variability of solar energy at diverse spatiotemporal scales. For short-term horizons up to an hour, sky cameras have demonstrated a unique capability to monitor the surrounding cloud cover, and thus predict the future impact of its dynamic on the local energy production. In this study, the generalisation skills of solar forecasting algorithms using neural networks to process cloud cover observations is evaluated beyond their training location via transfer learning. More specifically, two data scarce conditions are investigated: zero-shot learning (i.e. direct application of a trained model to a new location without local fine-tuning) and few-shot learning (i.e. calibration of a pre-trained model based on very limited local data such as a day of observations). Results demonstrate the potential of both methods to provide reliable predictions competitive with the ones of local experts trained on years of data. The use of a local clear-sky model significantly benefits ZSL methods by facilitating multi-domain (solar irradiance *versus* solar power production) target distribution matching; and FSL methods when local overcast sky images are used for model calibration. However, the generalisation to clear-sky conditions remains worse than that of the smart persistence model, partly because of the difficulty to design an accurate clear-sky model for solar power plants taking into account the spatial distribution of solar panels, their characteristics, and potential shading effects. In this context, fitting a clear-sky model with limited local data or applying bias correction methods could improve ZSL and FSL approaches further.

Overall, these results highlight the potential to apply computer vision-based solar forecasting pre-trained models to new locations with no or limited local data for model calibration. The increasing availability of large public and private sky image datasets collected around the world [53] is likely going to improve the technique further by enabling the development of so-called foundation models — large-scale models pretrained on di-

verse datasets and tasks [73] — that could easily be adapted to a new location of interest.

Acknowledgements The authors acknowledge SIRTA for providing the data used in this study. This research was supported by the Φ -Lab at the European Space Research Institute (ESRIN) and the Climate Office at the European Centre for Space Applications and Telecommunications (ECSAT) of the European Space Agency (ESA), EPSRC (EP/R513180/1), the University of Cambridge, and O.I.E. Centre Observation, Impacts, Energy at Mines Paris PSL. The authors would also like to thank the support by the Dubai Electricity and Water Authority (DEWA), United Arab Emirates through their membership in the Stanford Energy Corporate Affiliates (SECA) program and the Michael Hammer Postdoctoral Fellowship from The Institute for Data, Systems, and Society (IDSS) at Massachusetts Institute of Technology (MIT).

References

- [1] L. Dubus, S. Muralidharan, A. Troccoli, What Does the Energy Industry Require from Meteorology?, in: A. Troccoli (Ed.), *Weather & Climate Services for the Energy Industry*, Springer International Publishing, Cham, 2018, pp. 41–63. doi:10.1007/978-3-319-68418-5.
- [2] International Energy Agency, *Power Systems in Transition: Challenges and Opportunities Ahead for Electricity Security*, OECD, 2020. doi:10.1787/4ad57c0e-en.
- [3] G. Boyle, *Renewable Electricity and the Grid: The Challenge of Variability*, Routledge, 2009.
- [4] R. Bessa, C. Moreira, B. Silva, M. Matos, Handling renewable energy variability and uncertainty in power systems operation, *WIREs Energy and Environment* 3 (2014) 156–178. doi:10.1002/wene.76.
- [5] L. E. Jones, *Renewable Energy Integration: Practical Management of Variability, Uncertainty, and Flexibility in Power Grids*, Academic Press, 2017.
- [6] S. Koochi-Fayegh, M. A. Rosen, A review of energy storage types, applications and recent developments, *Journal of Energy Storage* 27 (2020) 101047. doi:10.1016/j.est.2019.101047.
- [7] M. S. Javadi, M. Lotfi, M. Gough, A. E. Nezhad, S. F. Santos, J. P. S. Catalão, Optimal Spinning Reserve Allocation in Presence of Electrical Storage and Renewable Energy Sources, in: *2019 IEEE International Conference on Environment and Electrical Engineering and 2019 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*, 2019, pp. 1–6. doi:10.1109/EEEIC.2019.8783696.
- [8] I. Antonopoulos, V. Robu, B. Couraud, D. Kirli, S. Norbu, A. Kiprakis, D. Flynn, S. Elizondo-Gonzalez, S. Wattam, Artificial intelligence and machine learning approaches to energy demand-side response: A systematic review, *Renewable and Sustainable Energy Reviews* 130 (2020) 109899. doi:10.1016/j.rser.2020.109899.
- [9] I. Ghalekhondabi, E. Ardjmand, G. R. Weckman, W. A. Young, An overview of energy demand forecasting methods published in 2005–2015, *Energy Systems* 8 (2017) 411–447. doi:10.1007/s12667-016-0203-y.
- [10] G. Kariniotakis, *Renewable Energy Forecasting: From Models to Applications*, Woodhead Publishing, 2017.
- [11] M. A. Islam, H. S. Che, M. Hasanuzzaman, N. A. Rahim, Chapter 5 - Energy demand forecasting, in: MD. Hasanuzzaman, N. A. Rahim (Eds.), *Energy for Sustainable Development*, Academic Press, 2020, pp. 105–123. doi:10.1016/B978-0-12-814645-3.00005-5.
- [12] C. Sweeney, R. J. Bessa, J. Browell, P. Pinson, The future of forecasting for renewable energy, *WIREs Energy and Environment* 9 (2020) e365. doi:10.1002/wene.365.
- [13] D. Yang, W. Wang, C. A. Gueymard, T. Hong, J. Kleissl, J. Huang, M. J. Perez, R. Perez, J. M. Bright, X. Xia, D. van der Meer, I. M. Peters, A review of solar forecasting, its dependence on atmospheric sciences and implications for grid integration: Towards carbon neutrality, *Renewable and Sustainable Energy Reviews* 161 (2022) 112348. doi:10.1016/j.rser.2022.112348.
- [14] A. Berrada, K. Loudiyi, Operation, sizing, and economic evaluation of storage for solar and wind power plants, *Renewable and Sustainable Energy Reviews* 59 (2016) 1117–1129. doi:10.1016/j.rser.2016.01.048.
- [15] G. Nottton, M.-L. Nivet, C. Voyant, C. Paoli, C. Darras, F. Motte, A. Fouilloy, Intermittent and stochastic character of renewable energy sources: Consequences, cost of intermit-

- tence and benefit of forecasting, *Renewable and Sustainable Energy Reviews* 87 (2018) 96–105. doi:10.1016/j.rser.2018.02.007.
- [16] D. Kumar, H. D. Mathur, S. Bhanot, R. C. Bansal, Forecasting of solar and wind power using LSTM RNN for load frequency control in isolated microgrid, *International Journal of Modelling and Simulation* 41 (2021) 311–323. doi:10.1080/02286203.2020.1767840.
- [17] S. Kulakov, F. Ziel, The Impact of Renewable Energy Forecasts on Intraday Electricity Prices, *Economics of Energy & Environmental Policy* 10 (2021). doi:10.5547/2160-5890.10.1.skul. arXiv:1903.09641.
- [18] P. Blanc, J. Remund, L. Vallance, 6 - Short-term solar power forecasting based on satellite images, in: G. Kariniotakis (Ed.), *Renewable Energy Forecasting*, Woodhead Publishing Series in Energy, Woodhead Publishing, 2017, pp. 179–198. doi:10.1016/B978-0-08-100504-0.00006-8.
- [19] B. Juncklaus Martins, A. Cerentini, S. L. Mantelli Neto, T. Z. Loureiro Chaves, N. Moreira Branco, A. von Wangenheim, R. R  ther, J. Marian Arrais, Systematic Review of Nowcasting Approaches for Solar Energy Production based upon Ground-Based Cloud Imaging, *Solar Energy Advances* (2022) 100019. doi:10.1016/j.seja.2022.100019.
- [20] M. G. Schultz, C. Betancourt, B. Gong, F. Kleinert, M. Langguth, L. H. Leufen, A. Mozaffari, S. Stadler, Can deep learning beat numerical weather prediction?, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379 (2021) 20200097. doi:10.1098/rsta.2020.0097.
- [21] Q. Paletta, G. Terr  n-Serrano, Y. Nie, B. Li, J. Bieker, W. Zhang, L. Dubus, S. Dev, C. Feng, Advances in solar forecasting: Computer vision with deep learning, *Advances in Applied Energy* 11 (2023) 100150. doi:10.1016/j.adapen.2023.100150.
- [22] J. Zhang, R. Verschae, S. Nobuhara, J.-F. Lalonde, Deep photovoltaic nowcasting, *Solar Energy* 176 (2018) 267–276. doi:10.1016/j.solener.2018.10.024. arXiv:1810.06327.
- [23] H. Gao, M. Liu, Short-Term Solar Irradiance Prediction From Sky Images With a Clear Sky Model, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 2475–2483.
- [24] Q. Paletta, A. Hu, G. Arbod, J. Lasenby, ECLIPSE: Envisioning CCloud Induced Perturbations in Solar Energy, *Applied Energy* 326 (2022) 119924. doi:10.1016/j.apenergy.2022.119924.
- [25] A. H. Nielsen, A. Iosifidis, H. Karstoft, IrradianceNet: Spatiotemporal deep learning model for satellite-derived solar irradiance short-term forecasting, *Solar Energy* 228 (2021) 659–669. doi:10.1016/j.solener.2021.09.073.
- [26] F. Lin, Y. Zhang, J. Wang, Recent advances in intra-hour solar forecasting: A review of ground-based sky image methods, *International Journal of Forecasting* (2022). doi:10.1016/j.ijforecast.2021.11.002.
- [27] Q. Paletta, G. Arbod, J. Lasenby, Omnivision forecasting: Combining satellite and sky images for improved deterministic and probabilistic intra-hour solar energy predictions, *Applied Energy* 336 (2023) 120818. doi:10.1016/j.apenergy.2023.120818.
- [28] 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, *Solar Energy* 118 (2015) 496–519. doi:10.1016/j.solener.2015.05.037.
- [29] P. Kuhn, B. Nouri, S. Wilbert, C. Prah, N. Kozonek, T. Schmidt, Z. Yasser, L. Ramirez, L. Zarzalejo, A. Meyer, L. Vuilleumier, D. Heinemann, P. Blanc, R. Pitz-Paal, Validation of an all-sky imager-based nowcasting system for industrial PV plants, *Progress in Photovoltaics: Research and Applications* 26 (2018) 608–621. doi:10.1002/pip.2968.
- [30] B. Nouri, S. Wilbert, L. Segura, P. Kuhn, N. Hanrieder, A. Kazantzidis, T. Schmidt, L. Zarzalejo, P. Blanc, R. Pitz-Paal, Determination of cloud transmittance for all sky imager based solar nowcasting, *Solar Energy* 181 (2019) 251–263. doi:10.1016/j.solener.2019.02.004.
- [31] N. Straub, W. Herzberg, A. Dittmann, E. Lorenz, Blending of a novel all sky imager model with persistence and a satellite based model for high-resolution irradiance nowcasting, *Solar Energy* 269 (2024) 112319. doi:10.1016/j.solener.2024.112319.
- [32] C. Feng, J. Zhang, SolarNet: A Deep Convolutional Neural Network for Solar Forecasting via Sky Images, in: *2020 IEEE Power Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, 2020, pp. 1–5. doi:10.1109/ISGT45199.2020.9087703.
- [33] Q. Paletta, J. Lasenby, Convolutional Neural Networks Applied to Sky Images for Short-Term Solar Irradiance Forecasting, in: *EU PVSEC*, 2020, pp. 1834 – 1837. doi:10.4229/EUPVSEC20202020-6BV.5.15.
- [34] Y. Nie, Y. Sun, Y. Chen, R. Orsini, A. Brandt, PV power output prediction from sky images using convolutional neural network: The comparison of sky-condition-specific sub-models and an end-to-

- end model, *Journal of Renewable and Sustainable Energy* 12 (2020) 046101. doi:10.1063/5.0014016.
- [35] H. Wen, Y. Du, X. Chen, E. Lim, H. Wen, L. Jiang, W. Xiang, Deep Learning Based Multistep Solar Forecasting for PV Ramp-Rate Control Using Sky Images, *IEEE Transactions on Industrial Informatics* 17 (2021) 1397–1406. doi:10.1109/TII.2020.2987916.
- [36] G. Terrén-Serrano, M. Martínez-Ramon, Deep Learning for Intra-Hour Solar Forecasting with Fusion of Features Extracted from Infrared Sky Images, 2022. doi:10.2139/ssrn.4064631.
- [37] Y. Sun, V. Venugopal, A. R. Brandt, Short-term solar power forecast with deep learning: Exploring optimal input and output configuration, *Solar Energy* 188 (2019) 730–741. doi:10.1016/j.solener.2019.06.041.
- [38] T. A. Siddiqui, S. Bharadwaj, S. Kalyanaraman, A Deep Learning Approach to Solar-Irradiance Forecasting in Sky-Videos, in: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), 2019, pp. 2166–2174. doi:10.1109/WACV.2019.00234.
- [39] Q. Paletta, G. Arbod, J. Lasenby, Benchmarking of deep learning irradiance forecasting models from sky images – An in-depth analysis, *Solar Energy* 224 (2021) 855–867. doi:10.1016/j.solener.2021.05.056.
- [40] W. Kong, Y. Jia, Z. Y. Dong, K. Meng, S. Chai, Hybrid approaches based on deep whole-sky-image learning to photovoltaic generation forecasting, *Applied Energy* 280 (2020) 115875. doi:10.1016/j.apenergy.2020.115875.
- [41] S.-A. Logothetis, V. Salamalakis, B. Nouri, J. Remund, L. F. Zarzalejo, Y. Xie, S. Wilbert, E. Ntavelis, J. Nou, N. Hendrikx, L. Visser, M. Sengupta, M. Pó, R. Chauvin, S. Grieu, N. Blum, W. van Sark, A. Kazantzidis, Solar Irradiance Ramp Forecasting Based on All-Sky Imagers, *Energies* 15 (2022) 6191. doi:10.3390/en15176191.
- [42] G. Terrén-Serrano, M. Martínez-Ramon, Deep Learning for Intra-Hour Solar Forecasting with Fusion of Features Extracted from Infrared Sky Images, 2022. doi:10.2139/ssrn.4064631.
- [43] Y. Nie, E. Zelikman, A. Scott, Q. Paletta, A. Brandt, SkyGPT: Probabilistic Short-term Solar Forecasting Using Synthetic Sky Videos from Physics-constrained VideoGPT, 2023. doi:10.48550/arXiv.2306.11682. arXiv:2306.11682.
- [44] J. Qin, H. Jiang, N. Lu, L. Yao, C. Zhou, Enhancing solar PV output forecast by integrating ground and satellite observations with deep learning, *Renewable and Sustainable Energy Reviews* 167 (2022) 112680. doi:10.1016/j.rser.2022.112680.
- [45] A. S. Bansal, T. Bansal, D. Irwin, A moment in the sun: Solar nowcasting from multispectral satellite data using self-supervised learning, in: Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, E-Energy '22, Association for Computing Machinery, New York, NY, USA, 2022, pp. 251–262. doi:10.1145/3538637.3538854.
- [46] Z. Si, M. Yang, Y. Yu, T. Ding, M. Li, A Hybrid Photovoltaic Power Prediction Model Based on Multi-source Data Fusion and Deep Learning, in: 2020 IEEE 3rd Student Conference on Electrical Machines and Systems (SCEMS), 2020, pp. 608–613. doi:10.1109/SCEMS48876.2020.9352398.
- [47] H. Wang, Y. Liu, B. Zhou, C. Li, G. Cao, N. Voropai, E. Barakhtenko, Taxonomy research of artificial intelligence for deterministic solar power forecasting, *Energy Conversion and Management* 214 (2020) 112909. doi:10.1016/j.enconman.2020.112909.
- [48] F. Wang, Z. Zhang, H. Chai, Y. Yu, X. Lu, T. Wang, Y. Lin, Deep Learning Based Irradiance Mapping Model for Solar PV Power Forecasting Using Sky Image, in: 2019 IEEE Industry Applications Society Annual Meeting, 2019, pp. 1–9. doi:10.1109/IAS.2019.8912348.
- [49] M. Xiang, W. Cui, C. Wan, C. Zhao, A Sky Image-Based Hybrid Deep Learning Model for Nonparametric Probabilistic Forecasting of Solar Irradiance, in: 2021 International Conference on Power System Technology (POWERCON), 2021, pp. 946–952. doi:10.1109/POWERCON53785.2021.9697876.
- [50] C. Feng, W. Zhang, B.-M. Hodge, Y. Zhang, Occlusion-perturbed Deep Learning for Probabilistic Solar Forecasting via Sky Images, in: 2022 IEEE Power & Energy Society General Meeting (PESGM), 2022, pp. 1–5. doi:10.1109/PESGM48719.2022.9917222.
- [51] Q. Paletta, A. Hu, G. Arbod, P. Blanc, J. Lasenby, SPIN: Simplifying Polar Invariance for Neural networks Application to vision-based irradiance forecasting, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2022, pp. 5182–5191.
- [52] Y. Nie, A. S. Zamzam, A. Brandt, Resampling and data augmentation for short-term PV output prediction based on an imbalanced sky images dataset using convolutional neural networks, *Solar Energy* 224 (2021) 341–354. doi:10.1016/j.solener.2021.05.095.
- [53] Y. Nie, Q. Paletta, A. Scott, L. M. Pomares, G. Arbod, S. Sgouridis, J. Lasenby, A. Brandt, Sky-image-based solar forecasting using deep learning with multi-location data: Training models locally, globally or via transfer learning?, *Applied Energy* (2024).
- [54] K. R. Gorantla, A. Roy, Generalizable Solar Irradiation Prediction using Large Transformer Models

- with Sky Imagery, in: 2023 18th International Conference on Machine Vision and Applications (MVA), 2023, pp. 1–5. doi:10.23919/MVA57639.2023.10216081.
- [55] Y. Nie, X. Li, Q. Paletta, M. Aragon, A. Scott, A. Brandt, Open-source sky image datasets for solar forecasting with deep learning: A comprehensive survey, *Renewable and Sustainable Energy Reviews* 189 (2024) 113977. doi:10.1016/j.rser.2023.113977.
- [56] E. Sarmas, N. Dimitropoulos, V. Marinakis, Z. Mylona, H. Doukas, Transfer learning strategies for solar power forecasting under data scarcity, *Scientific Reports* 12 (2022) 14643. doi:10.1038/s41598-022-18516-x.
- [57] M. Haeffelin, L. Barthès, O. Bock, C. Boitel, S. Bony, D. Bouniol, H. Chepfer, M. Chiriaco, J. Cuesta, J. Delanoë, P. Drobinski, J.-L. Dufresne, C. Flamant, M. Grall, A. Hodzic, F. Hourdin, F. Lapouge, Y. Lemaître, A. Mathieu, Y. Morille, C. Naud, V. Noël, W. O’Hirok, J. Pelon, C. Pietras, A. Protat, B. Romand, G. Scialom, R. Vautard, SIRTa, a ground-based atmospheric observatory for cloud and aerosol research, *Annales Geophysicae* 23 (2005) 253–275. doi:10.5194/angeo-23-253-2005.
- [58] Y. Nie, X. Li, A. Scott, Y. Sun, V. Venugopal, A. Brandt, SKIPP’D: A SKy Images and Photovoltaic Power Generation Dataset for Short-term Solar Forecasting, 2022. doi:10.48550/arXiv.2207.00913. arXiv:arXiv:2207.00913.
- [59] A. Driemel, J. Augustine, K. Behrens, S. Colle, C. Cox, E. Cuevas-Agulló, F. M. Denn, T. Duprat, M. Fukuda, H. Grobe, M. Haeffelin, G. Hodges, N. Hyett, O. Ijima, A. Kallis, W. Knap, V. Kustov, C. N. Long, D. Longenecker, A. Lupi, M. Maturilli, M. Mimouni, L. Ntsangwane, H. Ogihara, X. Olano, M. Olegs, M. Omori, L. Passamani, E. B. Pereira, H. Schmithüsen, S. Schumacher, R. Sieger, J. Tamlyn, R. Vogt, L. Vuilleumier, X. Xia, A. Ohmura, G. König-Langlo, Baseline Surface Radiation Network (BSRN): Structure and data description (1992–2017), *Earth System Science Data* 10 (2018) 1491–1501. doi:10.5194/essd-10-1491-2018.
- [60] G. Parmar, R. Zhang, J.-Y. Zhu, On Buggy Resizing Libraries and Surprising Subtleties in FID Calculation, arXiv:2104.11222 [cs] (2021).
- [61] S. Cros, S. Briand, J. Badosa, Benchmarking different approaches to convert surface solar irradiance into pv power production : a case study with an operational forecast system for a roof-top pv farm, in: EMS Annual Meeting 2022, Bonn, Germany, 5–9 Sep 2022, 2022, p. 1. doi:10.5194/ems2022-576.
- [62] Y. Nie, X. Li, A. Scott, Y. Sun, V. Venugopal, A. Brandt, SKIPP’D: A SKy Images and Photovoltaic Power Generation Dataset for short-term solar forecasting, *Solar Energy* 255 (2023) 171–179. doi:10.1016/j.solener.2023.03.043.
- [63] N. A. Engerer, F. P. Mills, KPV: A clear-sky index for photovoltaics, *Solar Energy* 105 (2014) 679–693. doi:10.1016/j.solener.2014.04.019.
- [64] M. Lefèvre, A. Oumbe, P. Blanc, B. Espinar, B. Gschwind, Z. Qu, L. Wald, M. Schroedter-Homscheidt, C. Hoyer-Klick, A. Arola, A. Benedetti, J. W. Kaiser, J.-J. Morcrette, McClear: A new model estimating downwelling solar radiation at ground level in clear-sky conditions, *Atmospheric Measurement Techniques* 6 (2013) 2403–2418. doi:10.5194/amt-6-2403-2013.
- [65] A. da Rosa, *Fundamentals of Renewable Energy Processes*, 2009. doi:10.1016/B978-0-12-374639-9.X0001-2. arXiv:arXiv:1011.1669v3.
- [66] P. Bacher, H. Madsen, H. A. Nielsen, Online short-term solar power forecasting, *Solar Energy* 83 (2009) 1772–1783. doi:10.1016/j.solener.2009.05.016.
- [67] C. Feng, J. Zhang, SolarNet: A sky image-based deep convolutional neural network for intra-hour solar forecasting, *Solar Energy* 204 (2020) 71–78. doi:10.1016/j.solener.2020.03.083.
- [68] C. Feng, J. Zhang, W. Zhang, B.-M. Hodge, Convolutional neural networks for intra-hour solar forecasting based on sky image sequences, *Applied Energy* 310 (2022) 118438. doi:10.1016/j.apenergy.2021.118438.
- [69] D. Yang, S. Alessandrini, J. Antonanzas, F. Antonanzas-Torres, V. Badescu, H. G. Beyer, R. Blaga, J. Boland, J. M. Bright, C. F. M. Coimbra, M. David, A. Frimane, C. A. Gueymard, T. Hong, M. J. Kay, S. Killinger, J. Kleissl, P. Lauret, E. Lorenz, D. van der Meer, M. Paulescu, R. Perez, O. Perpiñán-Lamigueiro, I. M. Peters, G. Reikard, D. Renné, Y.-M. Saint-Drenan, Y. Shuai, R. Urraca, H. Verbois, F. Vignola, C. Voyant, J. Zhang, Verification of deterministic solar forecasts, *Solar Energy* 210 (2020) 20–37. doi:10.1016/j.solener.2020.04.019.
- [70] I. M. Insaf, H. M. K. D. Wickramathilaka, M. A. N. Upendra, G. M. R. I. Godaliyadda, M. P. B. Ekanayake, H. M. V. R. Herath, D. M. L. H. Dissawa, J. B. Ekanayake, Global Horizontal Irradiance Modeling from Sky Images Using ResNet Architectures, in: 2021 IEEE 16th International Conference on Industrial and Information Systems (ICIIS), 2021, pp. 239–244. doi:10.1109/ICIIS53135.2021.9660664.

- [71] T. M. Mercier, T. Rahman, A. Sabet, Solar Irradiance Anticipative Transformer, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 2064–2073.
- [72] Y. Hatanaka, Y. Glaser, G. Galgon, G. Torri, P. Sadowski, Diffusion Models for High-Resolution Solar Forecasts, 2023. doi:10.48550/arXiv.2302.00170. [arXiv:2302.00170](#).
- [73] Y. Nie, Q. Paletta, S. Wang, Skyimagenet: Towards a large-scale sky image dataset for solar power forecasting, ICLR 2024 Workshop on Tackling Climate Change with Machine Learning (2024).