Proc. of the International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME 2023)
19-21 July 2023, Tenerife, Canary Islands, Spain

Intra-hour Solar Irradiance Forecast in Multiple Locations using Deep Transfer Learning

Aditi Roy[1], Ulrich Muenz[1], Joachim Bamberger[1], Ti-chiun Chang[2], and Chris Bilby[3]
[1]Siemens Technology, [2]Merck, [3]STEM
{firstname.lastname}@siemens.com, tichiun@gmail.com, cbilby@gmail.com

***Abstract*—In recent years, solar power system installation imposes several challenges on the operations of local and regional power grids due to the inherent variability of ground-level solar irradiance. This work proposes a novel real-time solar forecast methodology for intra-hour solar irradiance based on deep transfer learning from ground-based sky imager for time horizons ranging from 5-15 min. There are three unique aspects of the proposed methodology: (1) a Deep Learning based algorithm development which is modeled as a classification approach rather than a traditional regression approach; (2) the use of the Transfer Learning technique to show generalization capability, robustness, and portability of baseline model in the newly deployed location where availability of enough data for training is typically scarce, and (3) redefinition of point-based irradiation forecast error estimation technique with a window-based one that is more intuitive and user-friendly.**

The system is developed using multiple years of irradiance and sky image recording in New Jersey and one-year data from Colorado, USA. The method is validated against ground telemetry from these two locations of diverse geographic and climatic conditions. Results show that the forecasting method proposed in this work is robust and highly accurate (8% MAPE error) for multiple locations deployment.

***Index Terms*—Deep Learning, Convolution Neural Networks, transfer learning, irradiation forecast**

I. Introduction

The spatial and temporal variability along with uncertainty in ground-level solar irradiance directly affects the stability and efficiency of a photovoltaic (PV) power generation station. The combined effect of atmospheric constituents such as water vapor, aerosols, and clouds causes such variability in solar irradiance [10]. Significant solar ramps might happen in power production due to the ground-level solar irradiation variation that raises various issues in distribution system management. Future irradiation drop event prediction is a cost-effective technology to develop a smart, dynamic power grid for grid regulation, load-following production, power scheduling, and unit commitment [9]. The research avenue of solar forecasting can be categorized into three types depending on the forecast horizon: intra-hour, intra-day (1-24 hours ahead), and day-ahead (1–3 days ahead).

This work is funded by the Department of Energy under grant number DE-EE0008769.

Ti-chiun Chang and Chris Bilby contributed to the work while working in Siemens Technology and Holly Cross Energy, respectively.

The proposed work focuses on intra-hour solar forecasting that is crucial for real-time grid balancing, unit commitment, storage system optimization, automatic generation control (AGC), and operating regulation reserves [21]. In intra-hour solar forecast, frequently happening solar ramp imposes a great challenge. To handle this issue, high temporal and spatial resolution hemispherical information of the cloud movement collected from local telemetry and ground-based sky imagery [19] are a popular choice [9]. Without cloud cover information, prediction from any data-driven methods faces the challenge of a time interval lag compared to ground truth ramps [5]. Since the accurate prediction of solar ramps is essential to solar integration applications, local-sensing methods are used to capture cloud cover information, to enhance the accuracy in predicting solar ramps. From the time series sky images, it is possible to obtain a good estimate of the cloud trajectory and thus predict when and how much the sunlight would be occluded by clouds advancing towards the sun [3], [25]. To obtain cloud trajectory, cloud segmentation, and motion estimation has to be performed.

In this work, we propose a novel Deep Learning based irradiation prediction algorithm that combines strong prior knowledge, statistical features, and data-driven learning for regression into a seamless framework. Given a sequence of sky images, where we know where the sun is through the calibration step, we first detect the presence of a cloud and its velocity to define a prediction zone. This prediction zone is where the clouds, if any, will likely move in to occlude the sun. Features about pixel values from the prediction zone, together with time stamps and cloud speed, are used to train multiple binary classifiers that will predict irradiation drop in the predefined interval. Results from multiple classifiers are combined in a unique way using a hierarchical decision tree to obtain the final prediction.

Our major achievement is in formulating the PV forecast framework, which is usually tackled by regression approaches, as a classification problem. The output of the model is the decision of whether irradiation will be dropped by some pre-defined amount in the pre-defined time horizon based on whether the clouds are present or not in the prediction zone. Moreover, we developed a Transfer Learning technique to show the generalization of the proposed approach for new location deployment when sufficient data from that location

979-8-3503-2297-2/23/31.00$2023 IEEE$

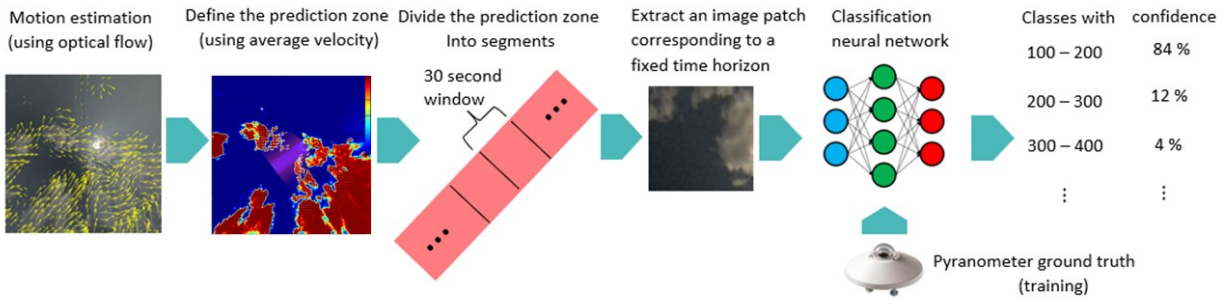

Fig. 1. The workflow of our classification-based solar prediction. Results of motion estimation (first image from left) and cloud detection (second image from left), with the prediction zone shown in the purple region extending from the sun position in the right image. The color of the segmentation map indicates the classification probability of the clouds, with dark red/blue being 1/0.

is not available to train any Deep Learning based PV forecast algorithm.

We tested the model on historical Princeton data and achieved an average accuracy of 92% for irradiation drop prediction events within a 5-15 minute time horizon irrespective of weather conditions. Next, we apply transfer learning on new location data from the Crystal substation of Holy Cross Energy (HCE) in Colorado, where the amount of training data was less than $\approx 30\%$ than that of the base location Princeton. Even with considerably less training data, the proposed approach was able to achieve the same level of forecasting accuracy. This study shows that a convolutional neural network (CNN)-based irradiation prediction algorithm with a classification approach offers great potential. The portability of the approach in the new locations using Transfer Learning with limited data also offers the possibility of developing such a system that can be used for new target regions in widely different climates.

II. Related Work

As mentioned before, to estimate sun occlusion and corresponding irradiation drop event, cloud trajectory estimation is an essential step. Cloud segmentation and motion computation are needed to predict cloud trajectories. Most earlier cloud segmentation techniques were based on color features [5], [11], [16]. Feature engineering is performed either by cloud detection methods [6], [11], [19] or by statistical RGB analysis methods [5]. For cloud detection, algorithms based on support vector machine (SVM) and Random Forest were used in the past [26]. Pedro et al. presented a detailed analysis of different sky image feature engineering methods [16], [17]. For motion estimation, Kalman filtering and variational optical flow-based methods were proposed in literature [4]. Most cloud-to-irradiance models for sky imagers are based on similar methods as satellite-based models, which are discussed in detail by Inman et al [9]. However, these image feature extraction methods are mostly manual in nature which increases the deployment and transfer costs. Conventional approaches, particularly for cloud detection, lack principled solutions to robustly adapt to the different sky (color) appearance from early morning to late afternoon, lens glares, and camera-induced small exposure change. To automatically and effectively obtain useful features from sky images, convolutional neural network (CNN)-based hybrid models are proposed in recent literature.

Deep learning-based models are commonly used nowadays for intra-hour solar forecasts [2] due to their ability to produce robust and generalizable irradiation prediction results [2], [20], [23]. CNN is one of the most established Artificial Neural Network structures [28], that surpassed the accuracy of classic segmentation techniques relying on color space thresholding. Other than CNNs, recurrent neural networks (RNNs), such as long short-term memory (LSTM) and gated recurrent unit (GRU) have been employed to predict irradiance as a time-series [1], [8], [23], [27], [29]. LSTM has been used to forecast short-term solar integrated load [18], [29], showing better accuracy against traditional models such as ARIMA, and multivariate linear regression. LSTM can "memorize" past information and make classification or prediction from new data that takes into account the intrinsic dynamics of the data. Since LSTMs require long training time, GRU has been applied for short-term PV forecasts [7], [8], [27].

The aforementioned approaches usually employ multiple sub-models to perform feature extractions and time series predictions separately. End-to-end methods that integrate multiple modules into a single framework have also been proposed to perform inference [13]. However, the end-to-end approach would have to incorporate motion estimation, cloud segmentation, and irradiance learning into a single black box, which could be risky to train without a careful design. So, we propose a multi-step solution where after detecting the presence of the cloud and its velocity separately, another module is used to predict irradiation drop in the predefined interval.

III. Methodology

Figure 1 shows the proposed workflow for classification-based solar irradiation prediction. Given 2 consecutive sky images, the first step is to segment the cloud pixels and estimate the cloud motion field using an optical flow technique. The yellow arrows in the left-most image show the motion vectors at the cloud locations. From the average cloud velocity, a prediction zone is defined and extracted as shown in red in the second image (from the left) in Fig. 1. Next features from the prediction zone are given as input to a Deep Neural Network (DNN) for irradiation drop prediction in the

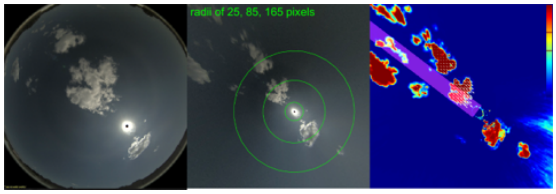

Fig. 2. Example results of motion estimation (left) and cloud detection (right), with the prediction zone shown in the purple region extending from the sun position in the right image. The color of the segmentation map indicates the classification probability of the clouds, with dark red/blue being 1/0.

predefined time horizon. Each of the steps is described in detail in the following subsections.

A. Cloud Analysis from Sky Images

To predict current to near-future sun occlusion, the presence of a cloud along with its motion should be estimated from the sky images. We model the first task as a binary pixel-wise classification (identifying cloud and non-cloud pixels), and the second task as an optical flow (i.e., 2D pixel motion map) regression. For cloud detection, we employed a neural network [15] which is trained using color and motion-based features. For motion estimation, we adopted dense optical flow implementation from OpenCV and found the results are reasonably good. Example results for cloud detection and motion estimation are shown in Fig. 2.

B. Prediction Zone Estimation from Global Cloud Motion

Next, we define a prediction zone in the sky images. If this zone contains any clouds, then there is a high probability that those clouds will move in to occlude the sun in the considered time horizon. We define the prediction zone as a 2D band [24] starting from the sun position and extending in the opposite direction of the global cloud motion (see a purple rectangle in Fig. 2), which is computed by averaging the motion orientations and speed.

First, the sun position s_i in 2D image space is obtained by projecting the current sun position in 3D sky coordinates obtained from "The Astronomical Almanac" [14], using the camera's extrinsic and intrinsic parameters. The global cloud motion vector is computed as $\overrightarrow{V_i} = \sum_{x,y=1}^{H,W} \gamma_i(x,y) \cdot v_{i \to i+1}(x,y)$ where $\gamma_i = h(m_i) \cdot f(\|d_i\|_2) \cdot g(v_{i \to i+1}, d_i)$, with $d_i(x,y) = (s_{i,x} - x, s_{i,y} - y)$ and h, f, and g functions of choices ($e.g.$, exponential or polynomial). The length of the prediction zone Z is proportional to the prediction time horizon and cloud speed. The orientation of Z is aligned with the incoming direction of the cloud's global average motion.

C. Binary DNN Classifier-based Irradiance Drop Prediction

Once the prediction zone is estimated, next it is segmented into multiple patches that correspond to different prediction time horizons, as illustrated in the third picture in Fig. 1. To make a prediction at the time horizon dT, a DNN takes as inputs a pair of patches that corresponds to time horizons 0 and dT. The output is a class label, which is encoded as an integer based on the irradiance difference between the current and

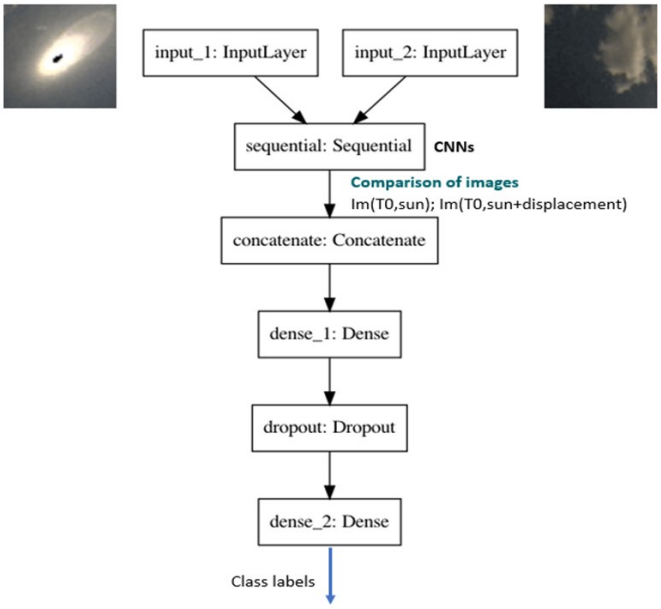

Fig. 3. Proposed DNN architecture that takes a pair of patches and outputs a class label. (b) One way to combine multiple binary classifiers for regression purposes.

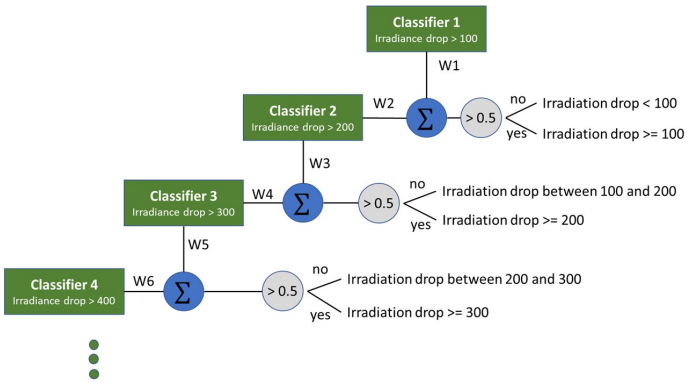

Fig. 4. Hierarchical Decision Tree to combine multiple binary classifiers.

future values. The right-hand side of Fig. 1 is further illustrated in Fig. 3, where the two extracted patches, an example DNN structure, and the output label are shown.

There can be two design choices for the DNN. The first is to have a multi-class classifier, with each class label corresponding to a predetermined value of irradiance drop. The second choice is to design multiple binary classifiers whose class labels correspond to a decision on whether the irradiance drop is larger than a predetermined threshold. The multiple binary classifiers can then be combined to make a prediction of different irradiance drops. We evaluated both approaches and identified that the predictions based on binary classification for a sequence of pairs provide better results.

Thus, to better approximate the amount of irradiance drop ranging from 100-1000 W/m^2, we developed multiple binary classifiers each one having class labels corresponding to a decision whether the irradiance drop is larger than a predetermined threshold. In this approach, we train one binary classi-

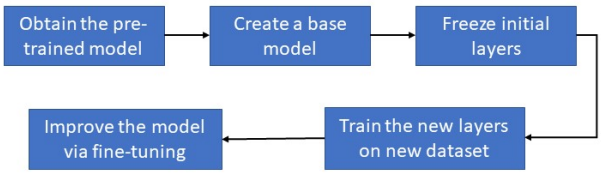

Fig. 5. Steps for the transfer learning process.

fier each for thresholds 100, 200, 300, 400, and 500 W/m^2, with each classifier predicting a drop at the corresponding threshold. Post training, we use a tree-based decision structure as shown in Fig. 4 to classify under which range the irradiance drop lies.

D. Transfer Learning Irradiation Prediction Network on New Location Data

The DNN model is initially trained using historical Princeton location data as sufficient multi-year data was available. However, the model has to be deployed in the Crystal substation of Holy Cross Energy (HCE) in Colorado location, where sufficient data was not available. For training another DNN model for the HCE location from scratch, it is not possible to get enough samples for a strong performance on the predictions. So, the trained model with Princeton data was used to predict on HCE data. As expected, a significant drop in the forecasting prediction accuracy (≈25% MAPE drop) was observed when trained and tested on data from different cameras and locations. Analysis of the results shows that the color appearance difference between training Princeton sky image data and test HCE sky image data played a major role in the accuracy drop. Additionally, cloud dynamics and camera field of view differences were other sources of errors. Such difference indicates the major divergence in the data domain that introduced bias in the model.

To address the above-mentioned issues of training with the new data set, the most popular solution is to apply Transfer Learning to "re-train" the DNN model with the new data available. Transfer learning leverages feature representations from a pre-trained model without requiring training a new model from scratch. The weights obtained from the pre-trained model are used to initialize the weights of the new model while training on new data. Including the pre-trained models in a new model leads to lower training time and lower generalization error. However, using pre-trained models for new tasks/domains is not that easy to apply as different data domains may have different feature spaces or different marginal distributions.

The steps followed for transfer learning are shown in Fig. 5. Once we get the Princeton pre-trained model, we instantiate the base model. The new model is initialized using the pre-trained weights. We kept the layers from the pre-trained model frozen to retain all the learning that has already taken place.

However, the base model has more units in the final output classification layer. When creating the base model, we explored both options of using the same classifier unit or adding a new output layer after removing the old complex one to suit our target problem. The second option worked best in our case where we settled for a single fully connected layer. Figure 6 demonstrates the modifications performed during transfer learning. Finally, we trained the model and performed finetuning to improve its performance. Fine-tuning is done by unfreezing the base model and training the entire model again on the whole new dataset at a very low learning rate. The low learning rate increases the performance of the model on the new HCE dataset while preventing over-fitting.

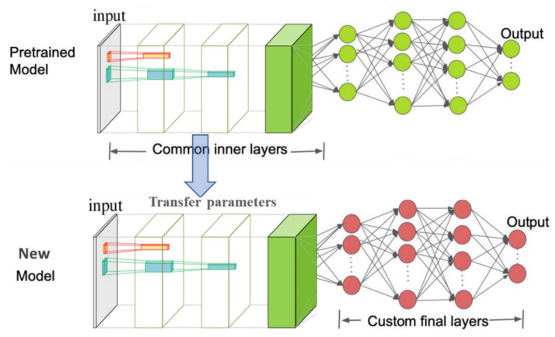

Fig. 6. Network architecture modification during transfer learning.

Different hyper-parameters were explored: (i) optimizers (e.g., stochastic gradient descent with and without momentum, ADAM, RMSprop), (ii) schedulers to adjust the learning rate based on the number of epochs (StepLR is used), (iii) regularization techniques to fight the overfitting problem. Further, the training dataset size was increased by sampling more drop events from the given time span raw data. Using this transfer learning approach, one binary classifier each for thresholds 100, 200, 300, 400, and 500 W/m^2 were trained. Post training, the decision structure shown in Fig. 4 was used to combine the decisions from multiple classifiers.

IV. Evaluation

A. Dataset

Data was collected in two locations: Princeton New Jersey, and Holy Cross Energy Crystal station (3.5MW) in Colorado, USA. The hemispheric sky images (180°/360°) were captured at an interval of 5 seconds spread over 12 hrs. per day with Mobotix MX-Q24 (Princeton location) and Mobotix Q26D (Colorado location) fish eye cameras. Corresponding irradiance measurements were collected using $Kipp\&Zonen$ SMP10 and SMP11 pyranometers in units of W/m^2 during the same time. The cameras were calibrated using the OcamCalib Toolbox [22] to obtain fisheye camera intrinsic parameters that are used for projecting the original image in Spherical coordinate to Cartesian coordinate for further image processing.

B. Evaluation Metric

We used Mean Absolute Percentage Error (MAPE) to measure prediction accuracy as used traditionally in statistics for evaluating a forecasting method [2], [12]. MAPE is the average of the absolute percentage errors of forecasts. Error is defined as the actual or observed value (d_t) minus the

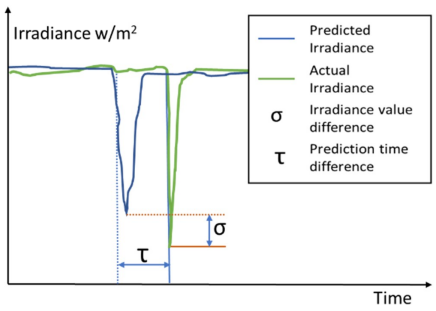

Fig. 7. Irradiation Prediction Evaluation Conditions.

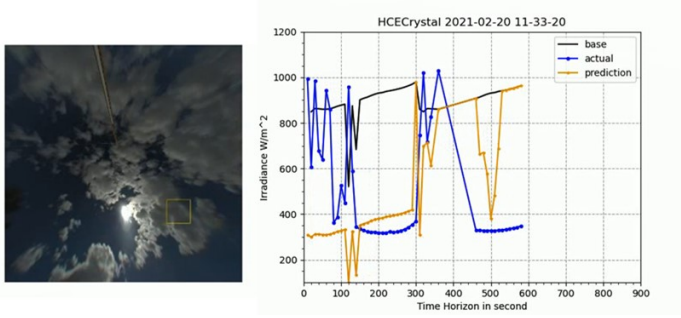

Fig. 8. Prediction for HCE data. Blue, yellow, and black curves are, respectively, the ground truth irradiance, the predictions, and the base values which we add the (predicted) differences to generate the final predictions. The orange boxes in the sky image are the patches used for classification for the 5-minute time horizon. Because the wind speed is high, not all prediction horizons are available in the images which have a limited field of view.

forecasted value(f_t). Percentage errors are summed without regard to sign to compute MAPE over all the n forecasted instances. It is defined as follows:

$$MAPE = \frac{100\%}{n} \sum_{t=1}^{n} |\frac{d_t - f_t}{d_t}| \quad (1)$$

For several applications, it is noticed that it would be sufficient if the predicted irradiation drop value is within $\pm\sigma$ of the actual irradiation drop value, where the actual value of σ is user-defined (refer Fig. 7). Moreover, if the predicted drop event is within $\pm\tau$ second of the actual drop event, then it would not have a significant impact on the downstream application (refer Fig. 7). Therefore, during the evaluation of the proposed approaches, we computed MAPE for different values of σ and τ that show the robustness of our approach in varied conditions.

C. Results

As mentioned before in Section II.D, the proposed DNN model was initially trained using historical Princeton data captured during 2015 and 2017, which consist of 20,175 noise-free samples. We evaluate the forecasting algorithm using 5,000 samples of 2016 Princeton data (i.e., unseen data during training) on three different prediction horizons: 5, 10, and 15 minutes. Results of evaluation for different values of σ and τ are shown in Table 1. It can be observed that during a 5-min ahead prediction horizon, if the predicted irradiation drop value w.r.t the observed value is within $\pm 50 W/m^2$, the classifier can predict the event within $\pm 30 sec$ time window with an average 10% error. The accuracy increases if the allowable drop value difference or the event detection window is more relaxed. Our current results achieve an 8% average MAPE for 15 minutes ahead prediction horizon on Princeton test data. The proposed approach shows robustness against different parameter variations. Depending on user requirements, it offers flexibility in choosing the right parameters best suited for the application while performance is still guaranteed.

Next, we evaluated the performance of the model after transfer learning on HCE data was performed. One example prediction instance with frequent irradiance drop in a cloudy sky scenario is shown in Figure 8. Fine-tuning was done with 6,214 samples of HCE data and then evaluation was performed on 1584 samples of unseen data. Similar to Princeton data, performance evaluation was done using MAPE when the predicted irradiation drop value is within $\pm\sigma$ of the actual irradiation drop value, and the predicted drop event is within $\pm\tau$ second of the actual drop event (refer Fig. 7). Results of evaluation for different values of σ and τ are shown in Table 2. It can be observed that fine-tuning significantly improved overall performance. During the 5-min ahead prediction horizon, if the predicted irradiation drop value w.r.t the observed value is within $\pm 150 W/m^2$, the classifier can predict the event within $\pm 30 sec$ time window with an average 11% error. The accuracy increases if the allowable drop value difference or the event detection window is more relaxed. We achieved an average 8% MAPE within 1 minute of the actual drop event (within a 5-minute time horizon) which is quite acceptable considering the challenges, like, low training dataset size, and task complexity under high environmental variation. Thus, the proposed approach shows strong generalization capability, robustness, and portability of the baseline model in a newly deployed location where the availability of enough data for training is typically scarce.

V. Conclusion

In this paper, we propose novel a Deep Learning-based approach for short-term irradiation forecast based on a classification approach. This is a completely new direction to handle this problem which is usually tackled by a regression approach. Moreover, we developed a Transfer Learning technique to show the generalization capability of the proposed approach for new location deployment when sufficient data is not available to train any Deep Learning based PV forecast algorithm. Further, we redefined the evaluation criteria for irradiation time-series forecasting that is more intuitive and usable by users offering greater flexibility and higher accuracy.

On historical Princeton data, we obtained 92% average accuracy for irradiation drop prediction events within 1 minute of the actual drop event. Even after applying Transfer Learning on the Crystal substation of Holy Cross Energy data, we were able to achieve the same level of average irradiation prediction accuracy. This shows the robustness of the proposed approach while deployed under different locations with geographical, environmental, and sensor variations.

MAPE								
Prediction Horizon	$\tau = 30sec$				$\tau = 60sec$			
	$\sigma = 0$	$\sigma = 50$	$\sigma = 100$	$\sigma = 150$	$\sigma = 0$	$\sigma = 50$	$\sigma = 100$	$\sigma = 150$
5 min ahead prediction	12%	10%	9%	7%	8%	6%	5%	4%
10 min ahead prediction	12%	11%	9%	8%	9%	7%	6%	5%
15 min ahead prediction	17%	15%	14%	13%	12%	11%	10%	9%
Average	**11%**				**8%**			

TABLE I
PERFORMANCE OF THE PROPOSED DNN MODEL FOR MULTI-HORIZON IRRADIATION DROP PREDICTION TESTED ON 5K SAMPLES OF PRINCETON DATA.

MAPE								
Prediction Horizon	$\tau = 30sec$				$\tau = 60sec$			
	$\sigma = 0$	$\sigma = 50$	$\sigma = 100$	$\sigma = 150$	$\sigma = 0$	$\sigma = 50$	$\sigma = 100$	$\sigma = 150$
5 min ahead prediction (before fine-tuning)	23%	21%	21%	18%	21%	19%	18%	15%
5 min ahead prediction (after fine-tuning)	18%	15%	12%	10%	11%	9%	7%	5%
Average after fine-tuning	**14%**				**8%**			

TABLE II
PERFORMANCE OF DNN MODEL FOR 5-MINUTE AHEAD PREDICTION TESTED ON 1.5K HCE DATA SAMPLES.

REFERENCES

[1] A. Alzahrani, P. Shamsi, M. Ferdowsi, and C. Dagli. Solar irradiance forecasting using deep recurrent neural networks. In *2017 IEEE 6th international conference on renewable energy research and applications (ICRERA)*, pages 988–994. Ieee, 2017.

[2] D. Anagnostos, T. Schmidt, S. Cavadias, D. Soudris, J. Poortmans, and F. Catthoor. A method for detailed, short-term energy yield forecasting of photovoltaic installations. *Renewable Energy*, 130:122–129, 2019.

[3] T.-c. Chang, J. Ernst, J.-R. Wiles, J. Bamberger, and A. Szabo. Short term cloud forecast, improved cloud recognition and prediction and uncertainty index estimation, May 28 2019. US Patent 10,303,942.

[4] C. W. Chow, S. Belongie, and J. Kleissl. Cloud motion and stability estimation for intra-hour solar forecasting. *Solar Energy*, 115:645–655, 2015.

[5] Y. Chu, H. T. Pedro, M. Li, and C. F. Coimbra. Real-time forecasting of solar irradiance ramps with smart image processing. *Solar Energy*, 114:91–104, 2015.

[6] Y. Chu, H. T. Pedro, L. Nonnenmacher, R. H. Inman, Z. Liao, and C. F. Coimbra. A smart image-based cloud detection system for intrahour solar irradiance forecasts. *Journal of Atmospheric and Oceanic Technology*, 31(9):1995–2007, 2014.

[7] R. Dey and F. M. Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pages 1597–1600. IEEE, 2017.

[8] M. Hosseini, S. Katragadda, J. Wojtkiewicz, R. Gottumukkala, A. Maida, and T. L. Chambers. Direct normal irradiance forecasting using multivariate gated recurrent units. *Energies*, 13(15):3914, 2020.

[9] R. H. Inman, H. T. Pedro, and C. F. Coimbra. Solar forecasting methods for renewable energy integration. *Progress in energy and combustion science*, 39(6):535–576, 2013.

[10] M. Lave and J. Kleissl. Solar variability of four sites across the state of colorado. *Renewable Energy*, 35(12):2867–2873, 2010.

[11] R. Marquez and C. F. Coimbra. Intra-hour dni forecasting based on cloud tracking image analysis. *Solar Energy*, 91:327–336, 2013.

[12] Y. Nie, Y. Sun, Y. Chen, R. Orsini, and 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(4):046101, 2020.

[13] Y. Nie, A. S. Zamzam, and 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*, 224:341–354, 2021.

[14] U. N. Observatory. Astronomical almanac, 1985.

[15] R. B. Palm. Prediction as a candidate for learning deep hierarchical models of data. *Technical University of Denmark*, 5, 2012.

[16] H. T. Pedro, C. F. Coimbra, and P. Lauret. Adaptive image features for intra-hour solar forecasts. *Journal of Renewable and Sustainable Energy*, 11(3):036101, 2019.

[17] H. T. Pedro, D. P. Larson, and C. F. Coimbra. A comprehensive dataset for the accelerated development and benchmarking of solar forecasting methods. *Journal of Renewable and Sustainable Energy*, 11(3):036102, 2019.

[18] X. Qing and Y. Niu. Hourly day-ahead solar irradiance prediction using weather forecasts by lstm. *Energy*, 148:461–468, 2018.

[19] S. Quesada-Ruiz, Y. Chu, J. Tovar-Pescador, H. Pedro, and C. Coimbra. Cloud-tracking methodology for intra-hour dni forecasting. *Solar Energy*, 102:267–275, 2014.

[20] W. Richardson Jr, D. Cañadillas, A. Moncada, R. Guerrero-Lemus, L. Shephard, R. Vega-Avila, and H. Krishnaswami. Validation of all-sky imager technology and solar irradiance forecasting at three locations: Nrel, san antonio, texas, and the canary islands, spain. *Applied Sciences*, 9(4):684, 2019.

[21] S. Sayeef, S. Heslop, D. Cornforth, T. Moore, S. Percy, J. Ward, A. Berry, and D. Rowe. Solar intermittency: Australia's clean energy challenge. characterising the effect of high penetration solar intermittency on australian electricity networks. 2012.

[22] D. Scaramuzza, A. Martinelli, and R. Siegwart. A toolbox for easily calibrating omnidirectional cameras. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5695–5701. IEEE, 2006.

[23] R. Sethi and J. Kleissl. Comparison of short-term load forecasting techniques. In *2020 IEEE Conference on Technologies for Sustainability (SusTech)*, pages 1–6. IEEE, 2020.

[24] S. Sun, J. Ernst, J. Bamberger, and J. R. Wiles. Sun location prediction in image space with astronomical almanac-based calibration using ground based camera, Oct. 22 2015. US Patent App. 14/711,002.

[25] S. Sun, J. Ernst, A. Sapkota, E. Ritzhaupt-Kleissl, J. Wiles, J. Bamberger, and T. Chen. Short term cloud coverage prediction using ground based all sky imager. In *2014 IEEE international conference on smart grid communications (SmartGridComm)*, pages 121–126. IEEE, 2014.

[26] S. Sun, J. Ernst, A. Sapkota, E. Ritzhaupt-Kleissl, J. Wiles, J. Bamberger, and T. Chen. Short term cloud coverage prediction using ground based all sky imager. In *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 121–126, Nov. 2014.

[27] J. Wojtkiewicz, M. Hosseini, R. Gottumukkala, and T. L. Chambers. Hour-ahead solar irradiance forecasting using multivariate gated recurrent units. *Energies*, 12(21):4055, 2019.

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

[29] Y. Yu, J. Cao, and J. Zhu. An lstm short-term solar irradiance forecasting under complicated weather conditions. *IEEE Access*, 7:145651–145666, 2019.