POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Enhanced Convolutional Neural Network for solar radiation nowcasting: All-Sky camera infrared images embedded with exogeneous parameters

Laurea Magistrale in Energy Engineering - Ingegneria Energetica

Author: Paolo Cusa

Advisor: Prof. Emanuele Giovanni Carlo Ogliari

Co-advisor: Alfredo Nespoli

Academic year: 2021-2022

1. Introduction

The ever-increasing energy demand leads to a larger spread of power generation in which renewable sources play a key role. The acceleration towards a clean energy transition is being pushed forward by stronger policies and ambitious climate targets announced for COP26. In particular, solar photovoltaic (PV) is expected to account for 60% of the worldwide annual renewable capacity additions [1]. Thus, the need for an accurate prediction of solar radiation is becoming crucial for both grid connections and stand-alone networks due to the intermittent and unpredictable nature of the solar source.

2. Forecast horizon and useful techniques

The forecast horizon is the time duration between the actual and the time of prediction. It can vary from a few minutes to several hours depending on the set goals. This study focuses on very-short term forecasting, also known as nowcasting, to make predictions at three horizons: 5, 10 and 15 minutes. The evaluation of solar radiation at such short timescales provides benefits for power smoothing processes, monitoring of real-time electricity dispatch and PV storage control. Nowcasting has lagged behind in research compared to longer forecast horizons. Nowadays, state-of-the-art is moving towards new technologies that are being developed to increase forecast reliability. Specifically, Convolutional Neural Networks (CNN) and thermal infrared All Sky Imagers (ASI) are becoming useful tools to improve accuracy. Liandrat et al.'s study makes use of those techniques to optimize a hybrid PV-Diesel system [2]. The Authors evaluate the effects of 10-minute solar prediction, comparing it to a scenario without forecast, and demonstrating how forecasting can reduce the generators' fuel consumption by up to 17% for equal numbers of blackouts in the microgrid. To determine how effective predictions of the proposed model are, benchmark methods are used for comparison. In the forecasting field, the most popular one is persistence, which belongs to the larger class of naïve methods. Despite their simplicity, which leads both to low computational and time costs, they are surprisingly effective and considered sufficiently accurate for very-short term forecasting.

3. Motivations and research questions

The aim of this thesis is to provide a forecasting technique for the immediate future based on the promising technologies that are developing nowadays: CNNs and infrared ASI. This goal is achieved by proposing a methodology which aims to solve two typical drawbacks related to machine learning and solar forecasting: the high computational burden and the difficulties of forecasting methods to outperform persistence. Moreover, an innovative technique in which exogenous data are inserted as information directly in the images is proposed. Lastly, a new analysis of the performance of the forecasting models is introduced for the first time. Specifically, an evaluation of the errors based on the instants in which consistent variations of solar radiation occur was performed.

4. Framework

Artificial Intelligence is a new frontier in data analysis due to its efficiency in terms of results and time with respect to the human brain. In particular, Convolutional Neural Networks (CNN) are suitable for image processing. They are built with three main types of layers: convolutional, pooling and fully connected layers. In the first ones, neurons work as a filter applied to the input. In case all the neurons in a layer use the same filter, the output is a map of activations called feature map. Filter parameters are the weights that the CNN learns throughout the training to detect features. Pooling layers aim to reduce the dimensionality of the representation, reducing the number of parameters and computational burden. Lastly, fully-connected layers make the vectorization of the previous layer output and provide values for regression purposes.

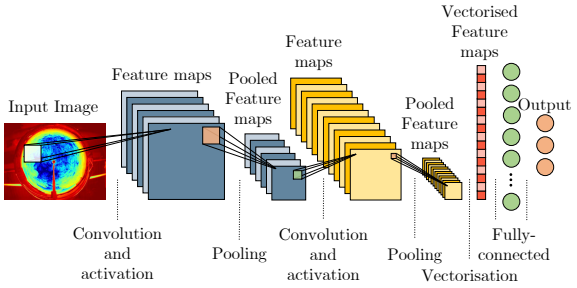

Figure 1: General architecture of a CNN

5. Proposed methodology

This proposed forecasting technique is based on a CNN which concatenates three images to create a sequence. The simultaneous visualization of consecutive images by the CNN model aims to detect clouds' motion which could affect irradiance in the following minutes. The proposed model can be represented as a non-linear mapping function that correlates the known inputs to the unknown output: the future Global Horizontal Irradiance (GHI). The problem is formulated as:

$$y_{\Delta t} = F_{\Delta t_{back}}(X, W) \tag{1}$$

where:

- $y_{\Delta t} \in \mathbb{R}^{N\times 1}$ is the future GHI;
- $X \in \mathbb{R}^{N\times W\times H\times C}$ is the sequence of images;
- N, W, H, C are the sample number, the image width and the image height, the number of channels of the image, respectively;
- Δt $\in \{5, 10, 15\}$ is the forecast horizon in terms of minutes;
- W is the trainable parameters matrix of the CNN;
- $\mathrm{F}_{\Delta \mathrm{t_{back}}}(\cdot)$ is the proposed CNN model whose inputs change and are trained accordingly with the *historical timestep* ($\Delta \mathrm{t_{back}}$);
- $\Delta \mathrm{t_{back}} \in \{2, 4, 8\}$ represents the time interval in terms of minutes between the images used to build the sequence.

5.1. Pre-processing

The available image data set was pre-processed with the aim of obtaining better results both in terms of performance and execution time. The first step was the cleaning phase, in which images depicting rainfall are removed. Precipitation is not the only phenomenon that can affect the forecast with infrared images. The thermal infrared all sky imager has a problem in case of low Sun elevation angle due to the increase of the atmospheric thickness (see Figure 2). Moreover, at Sunrise and Sunset, the region of the sky around the Sun is limited due to the presence of land. Hence, in this area it is particularly hard to detect clouds and their movements. This is the reason why images characterized by a Sun elevation angle lower than 20° were not considered by the forecasting model.

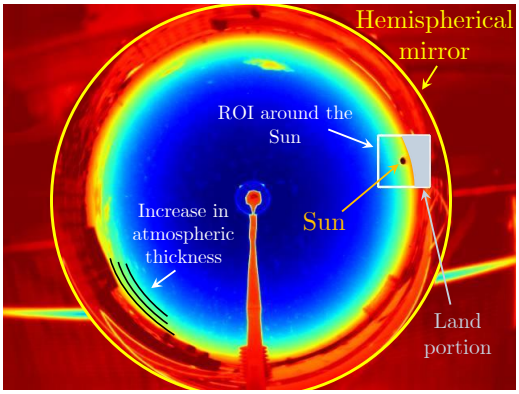

Figure 2: Infrared image and its characteristics

The second step is the conversion phase. Here, the goal is to convert the images so that they pass from Red-Green-Blue (RGB) format (three channels) to the greyscale format (one channel). This operation aims to accelerate the training process since the amount of image data is reduced by two-thirds. The image manipulation consists in making the inverse process of the one realized at the moment of the shot. In that case, a colormap was applied to a greyscale picture with a bit depth equal to 8, generating an RGB image with 256 shades. Jet function implemented in Matlab was used to create the same colormap from blue to red. Each RGB pixel value of the original image was compared with the 256 colors generated by the jet function. Among them, the one with the least difference from the RGB pixel was selected. Its index represents the grey intensity value for the monochrome image.

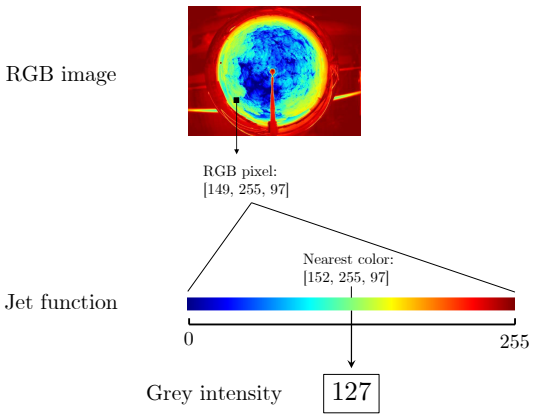

Figure 3: Illustration of the image transformation process from RGB to greyscale

The second phase is concluded through the application of a mask to remove all those parts which are not the sky. The last phase was cropping. Here, a square area of 384×384 pixels size was cropped from the original picture and scaling was performed to reduce the image size by two-thirds. In this way, unnecessary pixels were removed and a final image with 128×128 resolution was generated to reduce the data that has to be elaborated by the CNN.

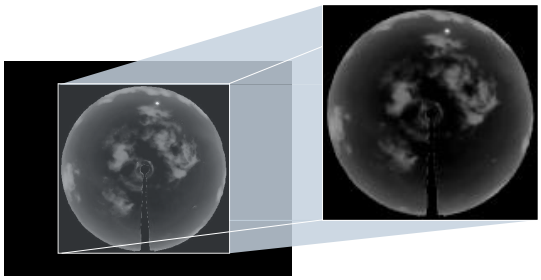

Figure 4: Cropping and scaling

5.2. Training

Image sequences are generated by concatenating them along the channel dimension. Hence, suppose that k images are concatenated as the input of the forecasting model, the input has the size of $H \times W \times (k \cdot C)$. The nowcasting task is to forecast the ramp events in the near future. Thus, sequences with last images taken before the prediction were concatenated. This operation is formally described in Equation 2.

$$I_{input} = (I_{t-(l-1)\Delta t_{back}} \oplus I_{t-(l-2)\Delta t_{back}} \oplus \cdots \oplus I_t) \tag{2}$$

In the equation, the symbol $\oplus$ indicates the concatenation between images (I), t is the actual time, l is the sequence length and Δt_{back} is the time interval between images in the sequence.
During training, the image sequences were passed multiple times through the CNN. Each iteration of the whole dataset corresponds to an epoch, at the end of which network parameters (the weights) are updated. The early stopping technique was used to stop the model training at the minimum of the error curve on the validation subset.
The clouds' motion extrapolated from the sequences is correlated with the global horizontal irradiance through the labels. They are associated with each image and constitute the solution for the CNN during the training phase. In

particular, the model does not predict directly the GHI, but the GHI normalized by the Clear Sky Global Horizontal Irradiance (GHI_{cs}). This last term was computed through the Ineichen and Perez clear sky model which was simulated in the pvlib Python package [3]. The resulting parameter is called Clear Sky Index (CSI) and it is a dimensionless parameter which directly informs on the weather conditions. Each sequence of images has to be associated with a label which was defined considering the forecast horizon. Thus, the target corresponds to the CSI value measured by waiting for the forecast horizon from the last image in the chronological order of the sequence. Considering a sequence composed of three images taken at t, t-Δt_{back} and t-$2\Delta t_{back}$, the CSI associated with that sequence is the one measured at t+Δt.

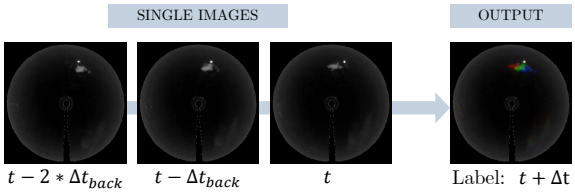

Figure 5: Sequence of images

5.3. Enhanced Convolutional Neural Network (ECNN)

The ECNN is a CNN to which exogenous data information is provided in the form of images. This model exploits a further step in image manipulation in order to improve the CNN based forecasting. Moreover, the ECNN was developed in order to maintain the same structure of the neural network and avoid the adoption of more complex forecasting methods which could handle different types of data. Exogenous data were codified into pixels and placed horizontally in the upper left corner of the image to not affect the sky region. Specifically, the data of interest selected was the solar radiation due to its strong relation with photos of the sky. In particular, the overall sum of the modified pixels is exactly the GHI, as shown in Figure 6. In this process, it was taken into account that the maximum value of an 8-bit pixel is 255 (2^8-1). This method ensures the presence of the three solar radiation measures in the three images sequence.

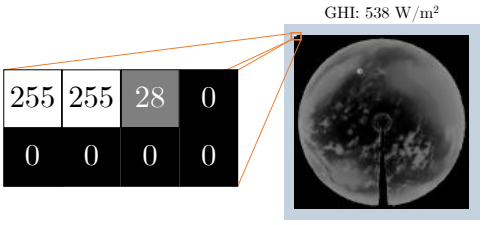

Figure 6: Codified exogenous data

5.4. Banchmarks and evaluation measures

Two models were used as a baseline for this work. The first one is the dataset of forecasts computed in 2020 by a weather broadcaster provider, which exploited, as the model proposed in this thesis, the infrared images and the GHI measurements collected by the meteorological station at the SolarTech^{LAB}. The second one is the persistence method, which basically assumes a constant CSI within the forecast horizon, resulting in a time-shifting of the GHI curve.
For the evaluation of the model, three indexes are selected: the *Root Mean Squared Error* (RMSE), the *Mean Bias error* (MBE) and the *Forecast Skill* (FS).

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_{\text{pred}} - y_{\text{meas}})^2} \quad (3)$$

$$MBE = \frac{1}{N}\sum_{i=1}^{N}(y_{\text{pred}} - y_{\text{meas}}) \quad (4)$$

$$FS = 1 - \frac{RMSE}{RMSE_p} \quad (5)$$

6. Case study

The dataset was collected from a thermal infrared all sky imager and a meteorological station both located in Milan, at the SolarTech^{LAB} of the Politecnico di Milano (latitude: 45.502921°N; longitude: 9.156564°E). Infrared all sky images and numerical meteorological measurements were collected in the period between 18th September 2019 and 9th April 2020 considering only the daylight hours, with an interruption period of recording between 23th December 2019 and 12th February 2020.
Images were taken with a resolution of 1 minute and their total number is equal to 116,891. After

filtering out rainy days and images where the Sun's elevation angle is too low, the available images used by the neural network are 33,686. They were shuffled by day in order to generalize the learning stage of CNN. Images were then concatenated to build the sequences and split into three subsets: *training*, *validation* and *test*; whose sizes are 64%, 16% and 20% respectively.

6.1. Thermal Infrared All Sky Imager

The Thermal Infrared All Sky Imager selected is the Sky InSight^{TM} developed by Reuniwatt. It includes a long-wavelength infrared camera that takes a picture with a 640×480 spatial resolution every 60 seconds. It is mounted on a mast to film a hemispherical mirror pointing at the sky, which provides a 180° field of view. Infrared radiation has broad prospects for providing valuable atmospheric properties. In fact, the spectral emission of sky which ranges from 8 to 13 μm is very sensitive to water vapor presence, facilitating the identification of clouds. The main advantages of infrared pictures compared to visible images are better overall cloud features as well as absence of Sun glare effect in Sun-region.

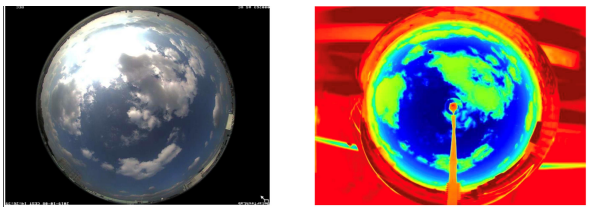

(a) (b)

Figure 7: Images taken in visible spectrum (7a) and infrared spectrum (7b)

However, a not negligible drawback is related to overcast weather. At those conditions, the clouds cover the entire sky and the infrared cam interprets as the sky dome having approximately the same temperature everywhere. This means that there is no possibility to distinguish sky and clouds as in the case of partly cloudy weather.

6.2. Meteorological station

The environmental conditions were collected with a weather station equipped with solar irradiance sensors, temperature and humidity sensors, a wind speed and direction sensor and a rain collector. Irradiance was measured with two secondary standard pyranometers, which evaluate the global irradiance on the horizontal and 30° tilted planes. Measurements were performed every ten seconds. The average, maximum, minimum and standard deviation of the values measured were calculated every minute.

6.3. CNN architecture

The architecture of the neural network is inspired by the one developed by the Visual Geometry Group (VGG) [4]. This structure is well-researched and widely used due to its good performance at image classification and low complexity. It is constituted by five Feature Learning Blocks (FLBs). Each FLB consists of two or three convolutional layers and a max-pooling layer. After the convolution part, there are two fully-connected layers. The first one has 256 neurons and the second has 1 neuron. This last neuron is responsible for providing the predicted value. The rectified linear activation function was used in convolutional layers, whereas the linear activation function was used in fully-connected layers. Lastly, a 0.2 dropout was put before each fully-connected layer to avoid overfitting. The model was trained with an initial learning rate of 0.0001 and the loss function adopted was the Mean Absolute Error.

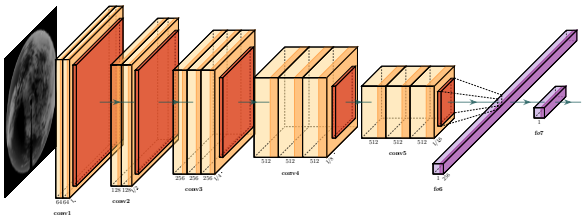

Figure 8: CNN architecture

7. Results

This study's results were obtained using an NVIDIA® T4 GPU on Google Colaboratory. The first outcome is expressed by the execution time. The average time of the simulations is less than two hours. This reveals the success in the reduction of computational costs.
The Root Mean Squared Error (RMSE), the Mean Bias Error (MBA) and the Forecast Skill (FS) were computed as functions of the forecast horizons and the historical timestep. Those metrics were then compared among the predictions provided by the weather broadcaster, per-

sistence and CNN method. The results show how the CNN model is able to outperform the predictions of the meteorological provider for all forecast horizons. Moreover, the proposed method matches the persistence method, except for predictions 5 minutes ahead due to the high accuracy of persistence at those timescales. The ECNN instead outperforms always persistence with FS of 5.05%, 9.57% and 9.07% for 5', 10' and 15'. Looking at the CNN models which got the best FS, they follow the real curve with good accuracy, but they are not able to outperform persistence. In particular, the CNN tends to behave like a very smart persistence model, avoiding large errors at the expense of missing peaks and having regular time delays. This behaviour is evident in Figure 9, in which all the ramps are missed, but both CNN and ECNN got FS higher than 14%. Lastly, Figure 10 shows the main problem of the proposed method: the overcast days. Here, the CNN models have errors at least three times larger than those observed with persistence. Those weather conditions are hard to evaluate by the CNN for two reasons: first, the impossibility of the CNN to know the Sun position; second, the similarity, in infrared images, between cloudy days and clear sky. The problem of overcast days is the main cause of the high positive bias in most of the analyses.

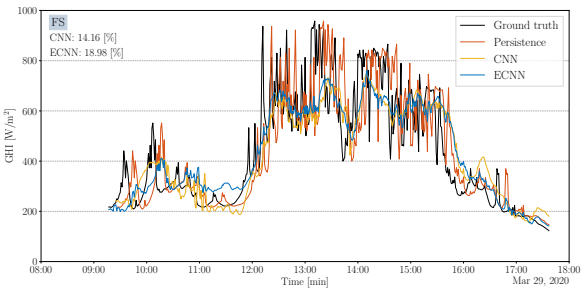

Figure 9: 29 March 2020 - partly cloudy

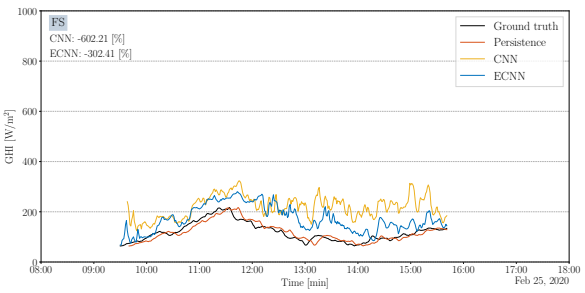

Figure 10: 25 February 2020 - overcast

7.1. Analysis on high fluctuations of irradiance

In literature, results provided by the forecasting models are usually carried out on the basis of different meteorological conditions. Specifically, they incur larger errors on partially cloudy days due to the greater variation of GHI compared to clear skies. In this thesis, a different approach is proposed. Results were directly analysed by looking at those instants characterized by high fluctuations in solar irradiance. Probability Density Functions (PDF) of the difference between the actual GHI and the future one were computed for each forecast horizon. Hence, three datasets for 5, 10 and 15 minutes ahead were obtained. The curve reported in Figure 11 shows that most of the irradiance variations are relatively small (lower than 60 W/m^2). However, the high fluctuations are fundamental for both electric grid balancing and microgrid management. The proposed analysis focuses precisely on this aspect, evaluating the performance of the best CNN models in terms of FS against the persistence method under those conditions that rarely occur, but which are significant. In particular, according to the bin width of each dataset, at least 80% of the most probable data were removed and not considered by the error metrics. Figure 11 depicts in light blue the region which was not evaluated, whereas red tails represent the analysed data.

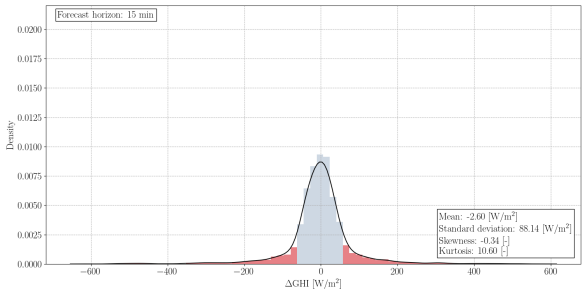

Figure 11: Probability density function for the 15 minutes forest horizon

The results show a significant improvement of the CNN models for these conditions, reaching forecast skill values always higher than 19% for all the forecast horizons. Those FS values demonstrate the effectiveness of this model for the most crucial and significant periods.

8. Conclusions

This study demonstrated the improvement that neural networks can bring to the forecasting field. However, this technique is still at an early stage of its development. Results highlight that persistence is still very accurate for very short-term forecasting (5 minutes) with respect to the basic CNN. However, the addition of exogenous data allows getting a significant improvement since ECNN always outperform persistence. Moreover, despite the positive FS (even improved in the case with high fluctuations of irradiance), the GHI curves reveal the tendency of the neural network to operate conservatively, favouring a smoothing trend in the predictions with respect to forecast ramps which are typical of partly cloudy days. The exploitation of larger data sets as well as training models on specific weather types could favour improvements.

9. Acknowledgements

I acknowledge Prof. Ogliari for his guidance and valuable advice. I also thank PhD Nespoli for his assistance and all the members of the energy department that helped me during this thesis.

References

[1] IEA. Renewables 2021 - analysis and forecast to 2026, 2021.

[2] Olivier Liandrat, Antonin Braun, Etienne Buessler, Marion Lafuma, Sylvain Cros, André Gomez, and Etienne Boudreault. Sky-imager forecasting for improved management of a hybrid photovoltaic-diesel system, 2018.

[3] William F Holmgren, Clifford W Hansen, and Mark A Mikofski. pvlib python: A python package for modeling solar energy systems. *Journal of Open Source Software*, 3(29):884, 2018.

[4] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.