Deep Learning for Solar Irradiance Nowcasting: A Comparison of a Recurrent Neural Network and Two Traditional Methods

Dennis Knol[1], Fons de Leeuw[2], Jan Fokke Meirink[3], and Valeria V. Krzhizhanovskaya[1,4]

[1] University of Amsterdam, Amsterdam, The Netherlands
[2] Dexter Energy Services, Amsterdam, The Netherlands
[3] Royal Netherlands Meteorological Institute (KNMI), De Bilt, The Netherlands
[4] ITMO University, Saint Petersburg, Russia

Abstract. This paper aims to improve short-term forecasting of clouds to accelerate the usability of solar energy. It compares the Convolutional Gated Recurrent Unit (ConvGRU) model to an optical flow baseline and the Numerical Weather Prediction (NWP) Weather Research and Forecast (WRF) model. The models are evaluated over 75 days in the summer of 2019 for an area covering the Netherlands, and it is studied under what circumstance the models perform best. The ConvGRU model proved to outperform both extrapolation-based methods and an operational NWP system in the precipitation domain. For our study, the model trains on sequences containing irradiance data from the Meteosat Second Generation Cloud Physical Properties (MSG-CPP) dataset. Additionally, we design an extension to the model, enabling the model also to exploit geographical data. The experimental results show that the ConvGRU outperforms the other methods in all weather conditions and improves the optical flow benchmark by 9% in terms of Mean Absolute Error (MAE). However, the ConvGRU prediction samples demonstrate that the model suffers from a blurry image problem, which causes cloud structures to smooth out over time. The optical flow model is better at representing cloud fields throughout the forecast. The WRF model performs best on clear days in terms of the Structural Similarity Index Metric (SSIM) but suffers from the simulation's short-range.

Keywords: Nowcasting · Solar Irradiance · Deep Learning · Convolutional GRU · Numerical Weather Prediction · WRF · Optical Flow.

1 Introduction

To help constrain global warming, the energy produced by fossil fuels is being replaced more and more by renewable energy. This transition is accelerated because wind and solar energy are now cheaper than energy from traditional resources [1,2]. However, in terms of usability, wind and solar are not yet fully competitive due to their variable nature. The way the energy market works dictates

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

that more accurate weather forecasts are essential to increase solar usability. The energy supply must continuously match demand, keeping the electricity grid in balance. As renewable resources replace conventional power plants, balancing the grid is becoming increasingly complex [3, 4].

One of the critical challenges is forecasting solar irradiance at the Earth's surface. When it comes to solar irradiance forecasting, clouds are the most critical driver and notoriously challenging to predict [5]. More traditional forecasting methods to forecast clouds divide roughly into two classes: image-based extrapolation methods (e.g., optical flow) and Numerical Weather Prediction (NWP) based methods. Extrapolation based methods perform well on a small temporal scale, but accuracy decreases markedly for increasing temporal scales as these methods do not take into account the evolution of clouds. NWP models generally have a better forecast accuracy for larger temporal scales, particularly for clear sky conditions [6]. However, performance decreases for cloudy conditions as NWP models do not explicitly resolve sub-grid cloud processes. Unresolved processes are parameterised and add a source of uncertainty to the model.

More recent studies propose to address these limitations from a machine learning perspective, taking advantage of the vast amount of weather data available. These studies formulate nowcasting as a spatiotemporal sequence forecasting problem and specifically focus on nowcasting precipitation, using Recurrent Neural Networks (RNN). Shi et al. [7] propose a novel Long Short Term Memory (LSTM) with convolutional layers to capture spatial correlation, which outperforms the optical flow based ROVER algorithm. A follow-up study proposes a Convolutional Gated Recurrent Unit (ConvGRU) [8], an architecture which is less complicated and more efficient to train. Google's MetNet is the first deep learning model to produce more accurate precipitation nowcasts than NWP [9].

This paper describes the implementation of a ConvGRU model for the nowcasting of solar irradiance trained on satellite observations. The ConvGRU model's performance is evaluated using two more traditional models as the benchmark: an optical flow algorithm and the Weather and Research Forecast (WRF) NWP model. We implement parameterisations specifically designed to meet the growing demand for specialised numerical forecast products for solar power applications for the latter. We also study the strengths and weaknesses of each method. A method's performance depends on the weather conditions (e.g., clear sky vs cloudy) and different models perform well on different temporal scales. This study aims to assess under what particular circumstances and spatiotemporal scales the models perform best.

2 Methods

2.1 Deep Neural Networks in the Nowcasting Domain

One of the most recent advances in the space of nowcasting is the use of Deep Neural Networks (DNN). Contrary to Optical Flow and NWP, deep neural networks can exploit the large amount of data collected continuously from ground-based cameras, radars, weather stations, and satellites.

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

Most progress is made in nowcasting precipitation. One of the early studies tackling the nowcasting problem from a deep learning perspective is [7]. In this study, researchers introduce an end-to-end trainable model, where both input and output are spatiotemporal sequences. The proposed model architecture is a Convolutional Long Short-Term Memory (ConvLSTM), which captures temporal dependencies using the LSTM cells and spatial dependencies with the convolutional layers. The model is trained using 812 days of radar data and outperformed a state-of-the-art operational optical flow algorithm called ROVER.

The follow-up study [8], led to the emergence of two more accurate model architectures: Convolutional Gated Recurrent Unit (ConvGRU) and Trajectory Gated Recurrent Unit (TrajGRU). The first updates the ConvLSTM by replacing the LSTM cells with Gated Recurrent Unit (GRU) cells. The latter differs from ConvLSTM and ConvGRU. In the TrajGRU model, convolutions generate flow fields, and this enables the model to learn the location-variant structures and capture spatiotemporal correlations more efficiently.

More recently, Google's researchers introduced MetNet and showed that the model outperforms the current operational NWP by the National Oceanic and Atmospheric Administration (NOAA), High-Resolution Rapid Refresh (HRRR), at predictions up to 8 hours. MetNet is a neural network model trained on both radar and satellite data. The data are first processed by a convolutional LSTM, second by axial attention layers. The attention mechanisms allow the model to ignore some parts of the data and focus on others and enable the model to learn long-term dependencies.

To our knowledge, solar nowcasting based on satellite data and deep learning techniques has not been covered by literature yet. The published literature on nowcasting irradiance with DNNs is based on sky-images retrieved from ground-based camera's [10,11]. This approach limits the forecast's lead time because the ground-based cameras cover only a small geographic area. Consequently, models trained on such data can only consider the clouds' possible motion over a very short period. Using satellite data, we can model larger geographical areas and generate forecasts further ahead in time.

2.2 Sequence-To-Sequence Model

Similar as in [7,8], this paper considers irradiance nowcasting a spatiotemporal sequence forecasting problem. A definition is presented in (1), where the input is a sequence of length J containing previous observations and the forecast a sequence of K frames ahead:

$$\tilde{I}_{t+1}, \ldots, \tilde{I}_{t+K} = \underset{I_{t+1},\ldots,I_{t+K}}{\arg\max}\ p\left(I_{t+1}, \ldots, I_{t+K} \mid \hat{I}_{t-J+1}, \hat{I}_{t-J+2}, \ldots, \hat{I}_t\right) \quad (1)$$

Here, the input and predictions are a sequence containing tensors, more formally $I \in R^{C\times H\times W}$. The spatial dimensions are H and W, the temporal dimension C. The input sequence contains the current observation, I_t. The model learns by minimising the forecast error through back-propagation and learns spatiotemporal relations without explicit assumptions about atmospheric physics.

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

For the problem, we adapt the encoder-decoder sequence-to-sequence structure proposed by [8]. This structure maps the input sequence with a fixed length on an output sequence with a fixed length and allows the input and output to have different lengths. An example is visualised in Figure 1. First, the encoder processes the elements in the spatiotemporal input sequence and creates a smaller and higher dimensional representation. Subsequently, the decoder learns to generate the predictions from the hidden state through multiple layers.

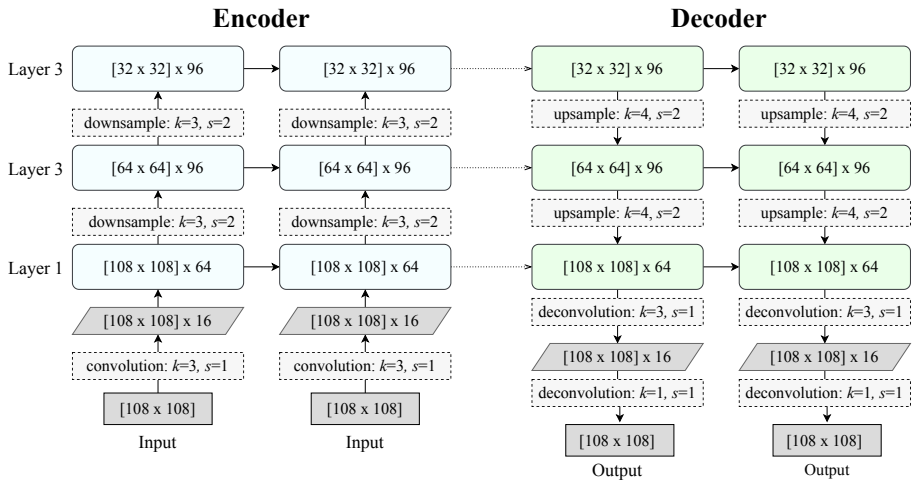

Fig. 1: Example of the sequence-to-sequence structure with three RNN layers, where $J = 2$ and $K = 2$. In the example, the model predicts the next two images based on the input of two images. Convolutions with stride, in between RNN layers, downsample the input in the encoder. Deconvolutional layers upsample the hidden representation in the decoder. We implement ConvGRUs as the RNN units, represented by the coloured cells. Figure inspired by [12].

2.3 Convolutional Gated Recurrent Unit

The ConvGRU convolutional layers assign importance to specific objects in the image, learn what parts of an input frame are important, and differentiate between different objects. The Gated Recurrent Unit (GRU), represented by the formulas in (2), captures the temporal dependencies by learning what previous information in a sequence is important to predict the future.

The GRU takes two sources of information: the memory state holding information of the previous units denoted with H_{t-1} and the information new to the model. The latter is denoted as X_t and is a single element of an inputted sequence. The first step in a GRU is to determine what information to pass on into the future using the update gate Z_t. Subsequently, the reset gate R_t determines what information to forget. With the output of the reset gate, the unit computes the new memory content H'_t through the activation function f. As the

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

last step, the unit computes the memory state H_t passed on to the next unit in the network. The corresponding formulas are as follows:

$$\begin{aligned} Z_t &= \sigma\left(W_{xz} * X_t + W_{hz} * H_{t-1}\right) \\ R_t &= \sigma\left(W_{xr} * X_t + W_{hr} * H_{t-1}\right) \\ H'_t &= f\left(W_{xh} * X_t + R_t \circ \left(W_{hh} * H_{t-1}\right)\right) \\ H_t &= \left(1 - Z_t\right) \circ H'_t + Z_t \circ H_{t-1} \end{aligned} \tag{2}$$

where $X_t \in \mathbb{R}^{C_i \times H \times W}$ and $H_t, R_t, Z_t, H'_t \in \mathbb{R}^{C_h \times H \times W}$. Here C is the number of input channels, H the height of the input image, and W the width. In the formulas, the convolutional operation is denoted as $*$, the Hadamard product as $\circ$ and f is the Leaky Rectified Linear Unit (ReLU) activation function with a negative slope of 0.2. The sigmoid activation function is applied to the update and reset gates.

2.4 Model Extension

Recent work in the medical field introduced model architectures that can exploit both static and dynamical medical data features, such as gender and patient visits [13]. The proposed architectures combine RNNs to process dynamic data with an independent Feed Forward Neural Networks processing the static data. The weather is also driven by a combination of dynamic and static features. Elevating terrain height, for example, causes orographic lifting and stimulates the formation of clouds. The ConvGRU model discussed above is specifically designed to work with sequence data and not exploit such data types. For the static input data, we develop an independent CNN with ReLU activation functions and batch normalisation, and we adapt the encoder-decoder structure. The encoder part of the network runs parallel to the encoder of the RNN and outputs the hidden state representation of the static information. Both encoders' hidden state is concatenated mid-way through the decoder and provide this information to the last layers, which output the irradiance prediction.

2.5 Optical Flow Baseline

The optical flow baseline used for this paper is an ensemble model of three optical flow algorithms available in the RainyMotion library [14]. The ensemble model, introduced by [15], combines the predictions of the Farneback, DeepFlow and Dense Inverse Search algorithms and provided more accurate irradiance forecast than the individual algorithms. The model is computed by taking the mean of the prediction of N optical flow algorithms:

$$\bar{e}_t(x, y) = \frac{1}{N} \sum_{i=1}^{N} p_t(x, y) \tag{3}$$

where i is the summation index, $\bar{e}_t(x, y)$ is the forecast of the ensemble model at lead time t and at coordinates x, y. The prediction of the optical flow algorithms is represented by p_t. We initialise the model on an input sequence of 5 frames containing the satellite-based irradiance parameter (see Section 2.7).

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

2.6 The WRF Baseline Model

Jimenez et al. [6] introduced an augmentation to the standard WRF model that makes the model appropriate for solar power forecasting. One of the studies that implemented the WRF model augmentation compares the model to the forecasts provided by the Global Forecasting System (GFS) and finds that the forecasting error is lower than of the GFS baseline [16]. However, this study does not consider Data Assimilation (DA), meaning that the WRF model suffers from incomplete information on the initial atmospheric state.

This limitation is addressed by a follow-up study [17]. This work incorporates satellite and ground-based observations into the WRF model's initial conditions using 3DVAR or 4DVAR initialisation. Both approaches improve the next-day forecast accuracy and the use of the latter results in the most accurate irradiance forecasts at all lead times. However, this comes at a computational expense. The 4DVAR initialisation takes 2 hours to compute on a high-performance machine. Given the steep increase in computational requirements and the studied forecasting window of 0 to 6 hours lead time, we did not apply DA in our study.

To set a baseline for this study, we use version 4.2 of the Advanced Research WRF. The 00:00UTC release of the Global Forecast System (GFS) provides the initial conditions of for the WRF model and updates the lateral boundary conditions hourly throughout the forecast. We set up the WRF model with a nested domain and implement the two-way nest option. The parent domain has a grid size of 60×60 and a horizontal resolution of 27 km, the middle point is at latitude 52.3702 and longitude 4.8952. The nested domain's resolution is 9 km and spans 70×70 grid cells, and we use the Lambert conformal conic projection. We summarise the important parameterisations in Table 1.

Table 1: The key physics settings used for the baseline WRF model.

Setting	**Implemented scheme**
Microphysics	Aerosol Aware Thomson micro-physics scheme [18]
SW Radiation	Rapid Radiative Transfer Model (RRTMG) scheme [19]
LW Radiation	Rapid Radiative Transfer Model (RRTMG) scheme [19]
Shallow Cumulus	Deng cumulus scheme [20]
Cumulus	Updated Kain-Fritsch cumulus scheme [21]
Land Surface model	Noah Land Surface Model [22]

2.7 Irradiance Data

We chose the irradiance data from the Meteosat Second Generation Cloud Physical Properties (MSG-CPP) algorithm developed by the KNMI to derive cloud parameters and solar radiation at the surface from the SEVIRI instrument onboard the Meteosat Second Generation satellite [23]. The data is available per 15 minutes, and the spatial resolution is 3×3 km^2. The MSG-CPP algorithm retrieves cloud properties, such as cloud optical thickness, thermodynamic phase and particle size, from the SEVIRI measurements of radiation reflected and emitted by the Earth and the atmosphere. Based on the computed cloud properties,

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

the algorithm derives the direct and diffuse surface irradiance components in W/m^2 [24]. We use the MSG-CPP irradiance data to initialise the optical flow and ConvGRU models and to evaluate the performance of all models.

A limitation of using this particular dataset is that the range of the processed data is somewhat shorter than the range of actual sunlight. The algorithm can only derive the cloud properties when the solar zenith angle (SZA) is not too high, as the estimations otherwise become very inaccurate. In the current version, the cloud properties are computed when the SZA is less than 78°, consequently missing a part of the day.

The data is retrieved from a geostationary satellite and spans an area much larger than the studied domain. For the preprocessing, we reproject the data to the Lambert Conformal Conic projection. Subsequently, we cut out the model domain presented in Section 3.

3 Experimental Setup

We next describe the overall setup for this research. We first compare different versions of the ConvGRU model. After that, we compare the best performing ConvGRU model to the baseline models. We consider a forecasting window of 0 to 6 hours and assess the forecasts at 15-minute intervals. We evaluate the models over 75 days in 2019, from July 18 until October 31.

This research's case study area is a geographical area corresponding to the Netherlands and is centred at latitude 52.4° and longitude 4.9°. For the WRF model, we rely on a nested domain setup. The optical flow and ConvGRU models' domain correspond to the outer domain in Figure 2. This domain provides a spatial context of approximately 300km in all direction to the domain of interest. The spatial context is essential as it provides the information to model incoming clouds. For all models, we evaluate the forecast in the inner domain.

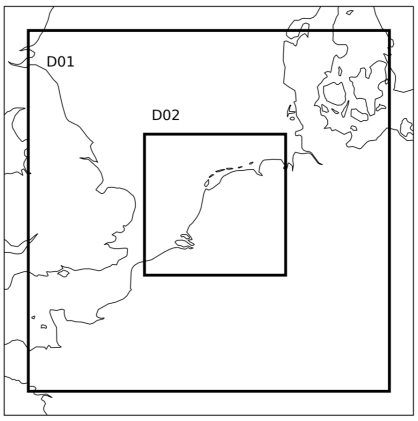

Fig. 2: Domain decomposition. The WRF domains have a spatial resolution of 9 and 27 km. The ConvGRU and optical flow models' resolution is 3km.

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

3.1 Training Set

The training data for the ConvGRU model comprises frames containing the MSG-CPP irradiance data, which have a resolution of 323×323 pixels. To reduce the computational requirements needed to train the models, we downsample the images to a resolution of 108×108. The ConvGRU model receives a four-dimensional tensor of size $[t, c, h, w]$ with dimensions time, number of channels, height and width. The time dimension of the input patch is 5, ranging from 08:00 - 09:00UTC as one frame is provided every 15 minutes. The output tensor contains 24 predicted frames, corresponding to 6 hours lead time. For training the models, we used 486 days of irradiance data (2018-03-15 - 2018-09-30, 2019-03-15 - 2019-07-17 and 2020-03-15, 2020-08-25) and the test set contains 75 days (2019-07-18 - 2019-09-30). By the time of the experiments, these were all the available days of MSG-CPP data with at least seven hours of irradiance data on the studied domain. We train all models using the Adam optimiser with a learning rate of 10^{-4}, 3000 epochs and a batch size of 6. The loss function we optimise is the smooth L1 loss function.

3.2 Forecast Evaluation

We calculate the Mean Bias Error (MBE) and the Mean Absolute Error (MAE) to determine the forecast accuracy of the models and quantify the strength of the error signal over the whole image. To complement the error-based metrics, we also compute the Structural Similarity Index Metric (SSIM) [25]. This metric accounts for patterns and textures in the images and can be used to assess the quality of the predicted images as it measures the loss of structural information by comparing local patterns in images. We compute the metrics based on normalised data to account for varying irradiance strengths throughout the day. We normalise the data with the clear sky irradiance from the MSG-CPP dataset.

Following [16], we assess the models' performance under different weather conditions. We compute Clear Sky Index (CSI) and categorise days into sunny, partly cloudy and very cloudy. We compute the CSI by dividing the average irradiance over a day by the average clear sky irradiance. We define days on which the average CSI is higher than 0.85 as sunny days and days with a lower CSI than 0.6 as mostly cloudy days. Days on which the CSI is between those values are defined as partly cloudy days.

4 Results

4.1 The ConvGRU Experiments

With the ConvGRU models, we conduct three experiments. The first experiment is based on the model introduced by [8] and trained on only a sequence of irradiance data. As we have a limited amount of data, we transpose the training data for the second experiment, doubling the number of images in the training set and increasing diversity. We thus train the ConvGRU model on twice as much

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

data. The third experiment is based on the model we introduced in Section 2.4. In this model, we input the sequence data into the ConvGRU and model the static data in the neural network specifically designed to model such features. The static features we add are terrain height, and land mask [26]. The frames containing these features cover the same spatial area as the WRF parent domain and have a 60 x 60 pixels resolution.

The model that performs best overall is the ConvGRU trained on only sequence data. Figure 3 shows an exemplary forecast, demonstrating that the model learned the advection of clouds. Additionally, we find that in all ConvGRU model experiments, image quality degrades rapidly and that structures are not preserved as a consequence of blurriness. The blurring effect, also visible in Figure 3, can be explained by the optimisation of a global-evaluating loss function and could be a result of the relatively small training set.

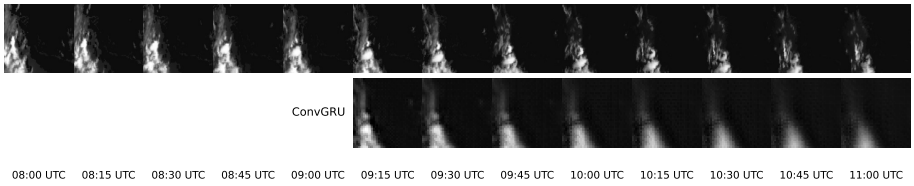

Fig. 3: Two-hour prediction sample of the ConvGRU model trained on only sequence data on 2019-07-24. The top row shows MSG-CPP data, of which the five left images are input to the model and the other images are the expected output. The second row shows the model's predictions. The normalised irradiance ranges from 0 (white) to 1 (black). One is the clear sky irradiance.

4.2 Comparison to the Baseline Models

Figure 4 contains prediction samples from each method. The optical flow algorithm generates the top prediction sample. In this prediction, clouds closely resemble the clouds in the expected output, and throughout the first two hours of the simulation, we note that the structures from the initialisation frames are preserved. This can be explained by the model's primary assumption that pixel values do not change over time. When comparing the methods, we note a clear difference at 09:15 UTC between the WRF prediction and the two others. This is because WRF is initialised on GFS data, while the ConvGRU and optical flow models are both initialised on the MSG-CPP data, and therefore, the output at the first lead time is very similar. After that, both generate increasingly different forecasts and portray the behaviour specific to both methods.

In the MSG-CPP images on the top rows of the figure, the shape of the clouds evolves over the forecast. In the optical flow predictions, the clouds found in the input sequence move across the images and are stretched out in the figure. The ConvGRU model blurs the images to optimise the global error metrics and as a result, causes cloud-like shapes to disappear over time. The WRF model is the

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

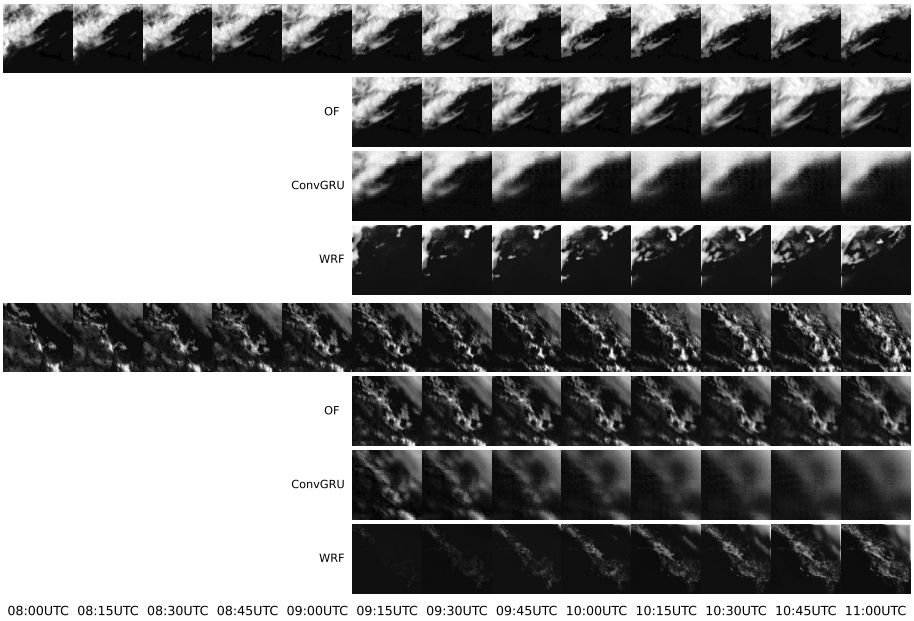

Fig. 4: Examples of two-hour irradiance predictions by the three different methods on 2019-07-24 (top) and 2019-08-20 (bottom). For each date, the images in the top row are the input and expected output, and the images in the second, third and fourth row are the predictions by the optical flow, ConvGRU and WRF model, respectively.

only forecasting method that can simulate the formation, growth and dissipation of clouds over time, but it suffers from inadequate initialisation and requires a spin-up time before more realistic cloud fields are simulated.

Figure 5 demonstrates that in terms of MBE and MAE, the WRF model is outperformed by the other models. In the first couple of model-steps, the optical flow and ConvGRU model performance are very similar, whereas the WRF

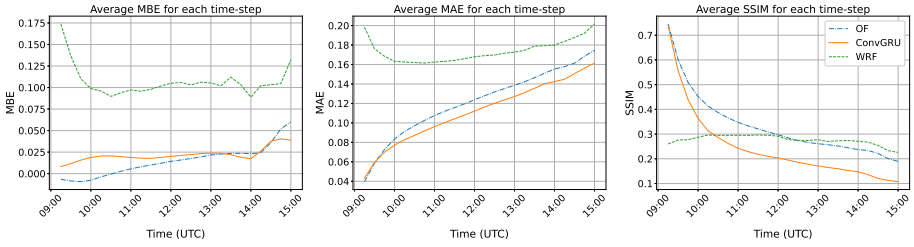

Fig. 5: The MBE, MAE and SSIM for the three methods. The metrics are averaged over the model domain and over all days from 2019-07-18 to 2019-09-30. Time is in UTC and the metrics are computed after normalisation of the data.

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

performance is particularly poor due to the spin-up problem. After about 45 minutes, ConvGRU outperforms optical flow in terms of MAE and, on average, improves the MAE by 9%. In terms of SSIM, the optical flow model is best during the first 3.5 hours, and the WRF model has the highest score after that. All models tend to overestimate irradiance (positive MBE) for nearly the entire forecast range in the study period.

4.3 Weather Dependent Analysis

We next examine the metrics in different weather conditions. In terms of MAE, it is clear that the ConvGRU outperforms the other methods in all weather conditions and from Figure 6, we see that the difference is largest on cloudy and very cloudy days. On such days, the ConvGRU model improves the optical flow MAE by 12% and 10%, respectively. On days with a clear sky index higher than 0.85, the improvement is 4%. The ConvGRU model, however, degrades more rapidly than the Optical Flow model in terms of SSIM. Especially on sunny days, optical flow is much better at predicting the irradiance than the ConvGRU model in terms of structural similarity. On average, the optical flow SSIM is 78% higher when compared to the ConvGRU SSIM.

Overall, the WRF model's performance is affected more strongly by the presence of clouds than the other models' performance. The MAE, for example, on partly cloudy days is more than double when compared to the same metrics for sunny days and increases by 42% for very cloudy days. Similarly, the MBE of the WRF forecast increases under more cloudy conditions (not shown), meaning that the model overestimates the irradiance more under such conditions. This demonstrates that the model consistently underestimates the irradiance absorbed and reflected (back to space) by clouds. The ConvGRU and optical flow predictions are less dependent on the weather than the WRF model.

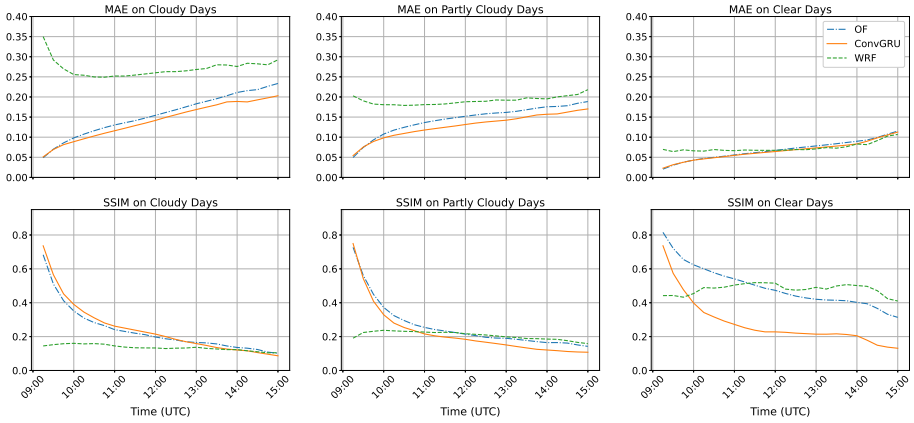

Fig. 6: The MAE (top) and SSIM (bottom) for all models grouped per weather type.

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

5 Conclusion

This research shows that the ConvGRU model trained on only sequences of MSG-CPP data provides the most accurate solar irradiance forecasts and improves the optical flow predictions by 9% in terms of MAE. The ConvGRU samples demonstrate that the model learned to represent clouds' advection, but it does show a blurring effect. The latter can be explained by the fact that the ConvGRU model is guided by a global-evaluating loss function and by the relatively small dataset the model is trained on. The blurring effect is the model's primary shortcoming as it causes the model not to represent cloud-like structures.

The cloud structures are better represented by the optical flow algorithm, especially on mostly clear days. The optical flow forecasts result in a 36% higher SSIM when compared to the ConvGRU predictions. Because of the assumption that pixel values stay constant from one frame to the next, optical flow mostly preserves the cloud-like structures from the initialisation data and advects the clouds over time. This enables the model to provide accurate predictions on a very short-term, but accuracy decreases for larger temporal scales as the optical flow model fails to represent the growth or dissipation of clouds.

One of the main strengths of trained DNNs and optical flow based methods is that models are easily initialised on the latest available weather data. Initialising WRF on the latest cloud data through data assimilation is a complex process, introduces uncertainty into the model and significantly increases the computational requirements. Because WRF was not initialised on the latest weather in our study, we cannot perform a reasonable comparison with the optical flow and ConvGRU models in the simulation's short range.

Which model is best depends on the forecast horizon and use cases. We also found that the performance of the different models depends on weather conditions. For that reason, operational forecasts are often provided by a forecasting system composed of several models. At lead times of 0 - 1 hour, the predictions by the optical flow model are best. On clear days and when cloud fields do not quickly dissipate or grow, the model performs well for a longer time. When compared to the ConvGRU model, optical flow is better at preserving the cloud field structure. This is a significant benefit when producing very short-term predictions of the PV output in a specific area, i.e. a solar farm. However, at longer lead times, predicting clearly defined cloud structures that are slightly different from the real clouds or in a different location can result in a very imprecise estimation of the future PV output of a specific solar farm. In such cases, the blurry output of the ConvGRU model might be preferred.

For future work, we recommend modifying the DNN model to optimise for a locally-oriented loss function to solve the blurry image problem of the ConvGRU model. Recent literature on precipitation nowcasting replaced the globally-oriented L1 loss function with a loss function based on the SSIM and showed that this approach improves the quality of the predictions and significantly reduces blurry image issue [12]. Furthermore, we recommend training the ConvGRU model on more data. This increases variability and generally improves the performance of RNNs and can be achieved in various ways.

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

References

1. Gimon, E., O'Boyle, M., Clack, C.T.: The coal cost crossover: economic viability of existing coal compared to new local wind and solar resources. Vibrant Clean Energy, LLC (2019)
2. Brunekreeft, G., Buchmann, M., Meyer, R.: New developments in electricity markets following large-scale integration of renewable energy. The Routledge Companion to Network Industries (2015)
3. Bird, L., Milligan, M., Lew, D.: Integrating Variable Renewable Energy: Challenges and Solutions. (2013)
4. Lund, P.D., Lindgren, J., Mikkola, J., Salpakari, J.: Review of energy system flexibility measures to enable high levels of variable renewable electricity. Renewable and Sustainable Energy Reviews. 45, 785–807 (2015)
5. Bauer, P., Thorpe, A., Brunet, G.: The quiet revolution of numerical weather prediction. Nature. 525, 47–55 (2015)
6. Jimenez, P.A., Hacker, J.P., Dudhia, J., Haupt, S.E., Ruiz-Arias, J.A., Gueymard, C.A., Thompson, G., Eidhammer, T., Deng, A.: WRF-Solar: Description and Clear-Sky Assessment of an Augmented NWP Model for Solar Power Prediction. 97, 1249–1264 (2016)
7. Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., Woo, W.: Convolutional LSTM network: A machine learning approach for precipitation nowcasting. arXiv preprint arXiv:1506.04214. (2015)
8. Shi, X., Gao, Z., Lausen, L., Wang, H., Yeung, D.-Y., Wong, W., Woo, W.: Deep learning for precipitation nowcasting: A benchmark and a new model. arXiv preprint arXiv:1706.03458. (2017)
9. Sonderby, C.K., Espeholt, L., Heek, J., Dehghani, M., Oliver, A., Salimans, T., Agrawal, S., Hickey, J., Kalchbrenner, N.: Metnet: A neural weather model for precipitation forecasting. arXiv preprint arXiv:2003.12140. (2020)
10. Siddiqui, T.A., Bharadwaj, S., Kalyanaraman, S.: A Deep Learning Approach to Solar-Irradiance Forecasting in Sky-Videos. 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). 2166–2174 (2019)
11. Zhang, J., Verschae, R., Nobuhara, S., Lalonde, J.-F.: Deep photovoltaic nowcasting. Solar Energy. 176, 267–276 (2018)
12. Tran, Q.-K., Song, S.: Computer Vision in Precipitation Nowcasting: Applying Image Quality Assessment Metrics for Training Deep Neural Networks. Atmosphere. 10, 244 (2019)
13. Esteban, C., Staeck, O., Baier, S., Yang, Y., Tresp, V.: Predicting Clinical Events by Combining Static and Dynamic Information Using Recurrent Neural Networks. IEEE International Conference on Healthcare Informatics (ICHI). (2016)
14. Ayzel, G., Heistermann, M., Winterrath, T.: Optical flow models as an open benchmark for radar-based precipitation nowcasting (rainymotion v0.1). Geoscientific Model Development. 12, 1387–1402 (2019)
15. Kellerhals, S.: Cloud and Solar Radiation Nowcasting using Optical Flow and Convolutional Gated Recurrent Units. MSc Thesis. University of Amsterdam. (2020)
16. Verbois, H., Huva, R., Rusydi, A., Walsh, W.: Solar irradiance forecasting in the tropics using numerical weather prediction and statistical learning. Solar Energy. 162, 265–277 (2018)
17. Huva, R., Verbois, H., Walsh, W.: Comparisons of next-day solar forecasting for Singapore using 3DVAR and 4DVAR data assimilation approaches with the WRF model. Renewable Energy. 147, 663–671 (2020)

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24

18. Thompson, G., Eidhammer, T.: A Study of Aerosol Impacts on Clouds and Precipitation Development in a Large Winter Cyclone. Journal of the atmospheric sciences. 71, 3636–3658 (2014)
19. Iacono, M.J., Delamere, J.S., Mlawer, E.J., Shephard, M.W., Clough, S.A., Collins, W.D.: Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models. Journal of Geophysical Research. 113, (2008)
20. Deng, A., Gaudet, B., Dudhia, J., Alapaty, K.: Implementation and evaluation of a new shallow convection scheme in WRF. In: 26th Conf. on Weather Analysis and Forecasting/22nd Conf. on Numerical Weather Prediction (2014)
21. Kain, J.S.: The Kain–Fritsch Convective Parameterization: An Update. Journal of applied meteorology. 43, 170–181 (2004)
22. Chen, F., Dudhia, J.: Coupling an Advanced Land Surface–Hydrology Model with the Penn State–NCAR MM5 Modeling System. Part I: Model Implementation and Sensitivity. Monthly weather review. 129, 569–585 (2001)
23. Roebeling, R.A., Feijt, A.J., Stammes, P.: Cloud property retrievals for climate monitoring: Implications of differences between Spinning Enhanced Visible and Infrared Imager (SEVIRI) on METEOSAT-8 and Advanced Very High Resolution Radiometer (AVHRR) on NOAA-17. Journal of Geophysical Research: Atmospheres. 111, (2006)
24. Greuell, W., Meirink, J.F., Wang, P.: Retrieval and validation of global, direct, and diffuse irradiance derived from SEVIRI satellite observations. Journal of Geophysical Research: Atmospheres. 118, 2340–2361 (2013)
25. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image Quality Assessment: From Error Visibility to Structural Similarity. IEEE transactions on image processing. 13, 600–612 (2004)
26. WRF Users Page: WPS V4 Geographical Static Data Downloads Page, `https://www2.mmm.ucar.edu/wrf/users/download/get_sources_wps_geog.html`, last accessed 2021/04/02.

ICCS Camera Ready Version 2021
To cite this paper please use the final published version:
DOI: 10.1007/978-3-030-77977-1_24