


APPLICATION PAPER

Week-ahead solar irradiance forecasting with deep sequence learning

Saumya Sinha^{1,*} , Bri-Mathias Hodge^{1,2} and Claire Montealeoni¹

¹Computer Science, University of Colorado, Boulder, Colorado, USA

²National Renewable Energy Laboratory (NREL), Boulder, Colorado, USA

*Corresponding author. E-mail: saumya.sinha@colorado.edu

Received: 28 September 2022; **Accepted:** 25 October 2022

Keywords: Deep sequence models; renewable energy; uncertainty estimation; week-ahead forecasting

Abstract

In order to enable widespread integration of solar energy into the power system, there is an increasing need to reduce the uncertainty associated with solar power output which requires major improvements in solar irradiance forecasting. While most recent works have addressed short-term (minutes or hours ahead) forecasting, through this work, we propose using deep sequence learning models for forecasting at longer lead times such as a week in advance, as this can play a significant role in future power system storage applications. Along with point forecasts, we also produce uncertainty estimates through probabilistic prediction and showcase the potential of our machine learning frameworks for a new and important application of longer lead time forecasting in this domain. Our study on the SURFRAD data over seven US cities compares various deep sequence models and the results are encouraging, demonstrating their superior performance against most benchmarks from the literature and a current machine learning-based probabilistic prediction baseline (previously applied to short-term solar forecasting).

Impact Statement

We show the promise of machine learning for longer-term solar forecasting with probabilistic predictions, an area that has not been sufficiently explored in the literature. Our encouraging results suggest such methods could play a larger role in future power system operations when greater shares of renewable energy resources will require operational planning at these timescales. For example, these methods could inform the operation of hybrid power plants with storage capabilities, where information about expected future renewable power generation would weigh into decisions on storage charging and discharging.

1. Introduction

Renewable energy sources, like solar, wind, tidal, or geothermal energy, have the potential of reducing the world's dependency on fossil fuel. These resources are not only abundantly available in nature but they are also clean energy sources, reducing greenhouse gas emissions that lead to global warming. However, many of these resources are variable and uncertain, posing challenges for integration into a power system which is predicated upon dispatchable supply. There is, therefore, a growing need for accurate renewable energy forecasting to ease integration into electric grids. Solar photovoltaics (PV) systems are experiencing exponential growth in deployment and the output of PV systems is highly dependent on solar irradiance (Alzahrani et al., 2017). A number of physical and statistical models have been used for making

solar forecasts at different timescales from intra-hour to a few days ahead (Tuohy et al., 2015; Wang et al., 2019). Statistical methods have been shown to perform well at forecasting at very short time horizons, with numerical weather prediction (NWP) models outperforming them in the hours to days-ahead timeframe (Tuohy et al., 2015).

Most physical models in this domain are based on NWP simulations that traditionally provide more accurate forecasts at hours to days-ahead lead times (Tuohy et al., 2015). However, due to their computational expense, NWP model outputs are updated less frequently and with coarser resolution at longer prediction lead times, such as week(s) ahead. This motivates the need for data-driven machine learning models that can provide forecasts at longer periods in advance at a finer (1 hr) resolution (as opposed to e.g., the 12 hr resolution in the case of the European Centre for Medium-Range Weather Forecasts (ECMWF) model predictions). As a part of our study, we not only perform a direct comparison with the NWP baseline for our 1-week ahead forecasts, but we also evaluate our models' performance when they incorporate NWP outputs as input features to see if it improves their forecasting ability.

Probabilistic forecasting provides a distribution over the prediction, this additional knowledge of uncertainty estimates can provide advantages over point forecasting. For example, knowing about future time periods of low and high uncertainty in advance can be very useful in planning plant maintenance (Zelikman et al., 2020). Until recently, probabilistic forecasting for solar energy had not received as much attention as for wind energy, as observed by Doubleday et al. (2020). In their work, Doubleday et al. (2020) introduce probabilistic benchmarks to evaluate probabilistic methods, which we will utilize in this work.

1.1. Contributions

We propose deep sequence learning for this longer lead-time (1 week ahead) solar irradiance forecasting task, that provides point as well as probabilistic predictions. Overall, these deep learning pipelines outperform several benchmarks from the literature including NWP models and a machine learning-based probabilistic prediction method. The results fall slightly behind the complete history persistence ensemble (Ch-PeEN) benchmark (Doubleday et al., 2020) in terms of continuous ranked probability score (CRPS), but are better in terms of forecast sharpness.

2. Related Work

A variety of deep learning approaches have been proposed for learning from sequence data, some of which have been applied in the solar energy domain. Recurrent neural networks (RNNs), unlike fully connected neural networks, have the ability to capture temporal dependencies in sequences by incorporating feedback from previous time steps. Long short-term memory (LSTM) models are especially useful for a time series data when the inputs can have longer dependencies. The works of Gensler et al. (2016), Alzahrani et al. (2017), Mishra and Palanisamy (2018), and Brahma and Wadhvani (2020) show the potential of LSTMs for solar energy forecasting, and they outperform fully connected networks and traditional machine learning models at short forecasting lead times. Convolution neural network (CNN)-based models that use dilated and causal convolutions along with residual connections (also referred to as temporal CNNs) were designed specifically for sequential modeling (Oord et al., 2016; Shaojie Bai and Koltun, 2018). They are autoregressive prediction models based on the recent WaveNet architecture (Oord et al., 2016). Temporal CNNs have recently been applied to forecasting day-ahead PV power output, outperforming both LSTMs and multilayer feed-forward networks (Lin et al., 2020). They are able to exploit a longer history in the time series, enabling more accurate forecasts. In this work, we study a significantly longer forecast horizon that challenges the limits of NWP forecasting and is expected to have emerging applications as power systems evolve. We compare LSTMs, temporal CNNs, temporal CNNs with an added attention layer (Saha et al., 2020), and the transformer model (Vaswani et al., 2017).

Recently, Zelikman et al. (2020) showed how probabilistic models such as Gaussian processes, neural networks with dropout for uncertainty estimation, and NGBoost (Duan et al., 2020) compare when

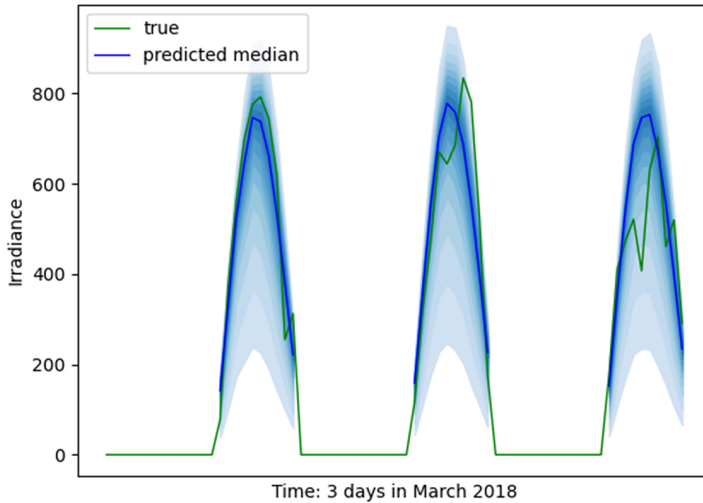


Figure 1. Fan plot showing the temporal CNN (TCN) model's prediction intervals from 5 to 95% percentile on three March days at the Boulder station.

making short-term solar forecasts. They explored post hoc calibration techniques for improving the forecasts produced by these models. NGBoost or Natural Gradient Boosting algorithm (Duan et al., 2020) is a gradient boosting pipeline that is extended to give a probabilistic distribution as an output (the parameters of the distribution are the regressed outputs). We now consider NGBoost with a Gaussian output distribution, to be a machine learning benchmark in this domain, since it showed superior performance for intra-hour and hourly resolution forecasting (Zelikman et al., 2020). Deep learning-based probabilistic prediction models are, however, yet to be fully explored (Wang et al., 2019). In this paper, we extend the deep learning point prediction models mentioned above to yield predictions at multiple quantiles (see Figure 1), as quantile regression is a nonparametric approach to obtain probabilistic forecasts (Wang et al., 2019; Saha et al., 2020).

3. Data

We use open-source NOAA's SURFRAD network (Surface radiation budget network for atmospheric research, Augustine et al., 2000) that provides the ground-truth solar irradiance and meteorological measurements from seven sites across the US in different climatic zones (<https://gml.noaa.gov/grad/surfrad/>). Models are trained on measurements from the years 2016–2017 and then evaluated in the year 2018. The test data (year 2018) is kept hidden and the rest of the data is split into training and validation sets (70/30 split). Data is converted to an hourly resolution and only the daytime values are considered for training and testing of all models including benchmarks (for relevance to the domain, as in Doubleday et al., 2020). Days with less than 24 hr of data points due to missing data were dropped. Following standard practice, we take a ratio of the ground-truth Global Horizontal Irradiance (GHI) (W/m^2) with respect to the “clear sky” GHI value (these are irradiance estimates under cloud-free conditions, obtained from CAMS McClear Service; Granier et al., 2019), to produce a clearness index, such as in Mishra and Palanisamy (2018), Doubleday et al. (2020), and Zelikman et al. (2020) that is used as the prediction label for training. While trained on the clearness index, the models are evaluated on the GHI.

Important predictor variables available in the data, such as solar zenith angle, hour of the day, month of the year, wind, pressure, temperature, and relative humidity, are included, along with the clearness index at the hour (a total of 16 input variables overall). These inputs are scaled (standardized) before the modeling procedure. All the sequence models take in a 3D input, where every row is a sequence of input feature vectors corresponding to previous timesteps: we use a history of 12×7 past daylight hours for all

our models. Each row in the time series at hour h is assigned a label which is the clearness index value at the hour $h + 1$ -week.

4. Methods

We focus on showing the potential of the following deep multivariate sequence models: LSTM, Temporal CNN, Temporal CNN with an attention layer, and Transformer, for point and probabilistic solar irradiance forecasting. We compare them to the NGBoost method (Duan et al., 2020) that has been shown to outperform various probabilistic models for short-term solar forecasting (Zelikman et al., 2020), along with benchmarks from the literature (as described in the next section). Hyperparameters were tuned on the validation dataset.

4.1. Long short-term memory

We use a simple LSTM pipeline; a single hidden layer with a dimension of 25.

4.2. Temporal CNN

Temporal CNN (TCN) consists of 1D dilated convolution filters and residual layers that are responsible for learning long-term dependencies efficiently (Shaojie Bai and Koltun, 2018; Lin et al., 2020). Figure 2 shows how dilations help to increase (exponentially) the receptive field of a kernel. This makes the model capable of learning correlations between data points far apart in the past. The convolutions are also causal, meaning that while convolving, outputs at time t only convolve with time t and earlier from the previous layer. Our TCN architecture is comprised of three levels, size of the hidden layer is 25, and kernel size is 3 with dilation factors $d = 1, 2$, and 4.

4.3. TCN with attention

The Attention mechanism (Cho and Bengio, 2015) has been used for sequential modeling and time series prediction problems (Qin et al., 2017). It has the ability to model dependencies in long sequences without regard to their distance (Vaswani et al., 2017). We add a self-attention layer (adapted from Zhang et al., 2019) on the convolution maps generated from the Temporal CNN network and observe the prediction outcomes. This enables the model to “pay attention” to various important parts of the feature maps that can help in making more accurate predictions.

4.4. Transformer

Transformers are architectures that are comprised of only the attention layers, leaving out any recurrence or convolutions entirely (Vaswani et al., 2017). They have been adapted for the task of time series forecasting as they work very well with longer sequences (Song et al., 2018; Wu et al., 2020). For this work, we use the encoder structure of transformers and work with a single stack of two-headed self-attention modules and other standard layers based on Vaswani et al. (2017).




Figure 2. Dilation in kernels (Oord et al., 2016; Borovykh et al., 2017).

4.5. Probabilistic prediction

For probabilistic forecasts, the above models are modified to output predictions at multiple quantiles (from 5 to 95%). While the point models are trained with mean squared error losses, their probabilistic counterparts are trained using quantile loss.

A fully connected layer at the end of each model is modified to produce either a single output (for point) or multiple outputs (for probabilistic). NGBoost is trained with default parameters and 2000 estimators as in Zelikman et al. (2020).

5. Evaluation

We provide the results of our experiments over all 7 SURFRAD stations for the test year (2018) in Tables 1 and 2. The benchmarks from the solar energy literature (derived from Doubleday et al., 2020) are:

5.1. Hourly climatology

Hourly climatology (HC) is a model that assigns the irradiance at a certain hour in 2018, to be the average of all irradiance values at the same hour of every day in the training data. For the probabilistic forecast evaluation, we do not use the average but the cumulative distribution function (CDF) over these values.

Table 1. Results of the point forecasting pipeline.

	SP	HC	NWP	NGBoost	LSTM	TCN	TCN + Attention	Transformer
Sioux Falls, SD	222.34	220.14	288.84	177.01	205.22	190.9	199.63	229.67
Fort Peck, MT	203.11	225.12	285.92	161.11	149.14	150.64	158.73	172.94
Bondville, IL	269.6	230.99	298.91	196.78	210.83	205.32	211.21	242.24
Penn State, PA	235.57	225.67	280.43	187.04	192.71	199.02	193.69	227.08
Boulder, CO	242.01	220.52	329.78	177.21	184.68	182.56	187.59	213.74
Desert Rock, NV	146.83	196.15	379.63	110.76	104.26	110.75	137.52	117.52
Goodwin Creek, MS	252.01	230.06	327.24	184.37	184.49	186.61	187.75	210.28

Note. Results are in terms of RMSE scores (lower the better). Comparisons are made with the SP, HC, and NWP benchmarks. Abbreviations: HC, hourly climatology; LSTM, long short-term memory; NWP, numerical weather prediction; RMSE, root mean squared error; SP, smart persistence; TCN, temporal CNN.

Table 2. Results of the probabilistic forecasting pipeline.

	HC	CH-PeEN	NWP	NGBoost	LSTM	TCN	TCN + Attention	Transformer
Sioux Falls, SD	126.04	87.75	218.42	106.41	93.96	94.03	98.14	153.29
Fort Peck, MT	129.37	77.37	217.76	97.83	80.05	76.65	77.83	95.59
Bondville, IL	131.79	100.84	225.15	119.07	117.79	107.87	110.68	134.31
Penn State, PA	126.9	97.74	199.22	114.38	101.37	107.43	103.3	131.04
Boulder, CO	126.46	88.65	256.79	102.91	92.42	91.47	91.25	114.77
Desert Rock, NV	112.46	44.78	311.33	55.13	45.21	46.43	56.43	51.11
Goodwin Creek, MS	130.4	95.91	249.16	111.05	95.26	98.6	100.63	120.91

Note. Results are in terms of CRPS scores. Comparisons are made with the probabilistic HC, CH-PeEN, and NWP benchmarks. The lower the CRPS, the better the model. Abbreviations: CH-PeEN, complete history persistence ensemble; CRPS, continuous ranked probability score; HC, hourly climatology; LSTM, long short-term memory; NWP, numerical weather prediction; TCN, temporal CNN.

5.2. Complete history persistence ensemble

CH-PeEN is used as a probabilistic prediction benchmark, where for a certain forecast hour, we take a CDF over the clearness indices at the same hour of every day from the training data and these are further converted to irradiance measures.

5.3. NWP ensembles

We are using the ECMWF 51-member ensemble as our NWP outputs. These members are updated only twice a day for 1-week ahead forecasts, and hence had to be repeated for the rest of the hours of the day (to be consistent with other forecasts). For point forecasts, we take the ensemble mean, while for probabilistic prediction we take an empirical CDF over them.

5.4. Smart persistence

Smart persistence (SP) is a model that assumes the clearness index (ratio of GHI/clear-sky GHI) at time t +lead-time to be the same as at time t , and uses that to obtain the irradiance at t + lead-time. This is a common benchmark from the short-term point forecasting literature, which we would not expect to perform well at longer forecast lead times, but include for the sake of completeness.

5.5. Evaluation metrics

The evaluation metrics for point forecasting are the RMSE (root mean squared error) scores of each model. For probabilistic forecasting, we use CRPS or Continuous Ranked Probability Score. CRPS is a widely used metric for evaluating probabilistic forecasts as it balances reliability, resolution, and sharpness which are other criteria to measure the quality of probabilistic outputs (Gneiting et al., 2007). Intuitively, CRPS measures the area between the predicted and the observed CDF, the observed (true) CDF being a step function at the observation (Zelikman et al., 2020). The lower the CRPS, the better the model. To evaluate our probabilistic forecasts, that is, when our model outputs predictions at different quantiles ($\zeta \in (0, 1)$), the CRPS score can be expressed as an integral over quantile scores(QS) at all quantiles (from Doubleday et al., 2020):

$$\text{CRPS} = \int_0^1 \frac{1}{T} \sum_{t=1}^T \text{QS}_{\zeta}(P^{-1}(\zeta, t), y(t)) d\zeta, \quad (1)$$

where y is the observation, 1 is an indicator function, P the predicted CDF distribution, and T the number of data points. QS at a particular ζ is defined as:

$$\text{QS}_{\zeta} = 2(1\{y(t) \leq P^{-1}(\zeta, t)\} - \zeta)(P^{-1}(\zeta, t) - y(t)). \quad (2)$$

Reliability looks at the statistical consistency between the forecast distribution and observed distribution, while sharpness looks at the concentration (narrowness) of the forecast (Gneiting et al., 2007; Lauret et al., 2019; Doubleday et al., 2020). These characteristics are best observed with a visual analysis and we follow the work of Doubleday et al. (2020) to visualize both reliability and sharpness. The sharpness plot in Figure 3b is where we plot the average forecast width at 10, 20, 30 %, ... central intervals. As we can see, sharpness does not look at the observation, it just considers the narrowness of the prediction interval. We also provide a reliability diagram (Figure 3b) where we compare the proportion of the observations that lie within a given quantile output, versus the quantile or nominal proportion itself (in an ideal scenario both are expected to be equal).

As we observe from Table 1, for the majority of the stations, all of our proposed deep learning models including LSTM, TCN, TCN + Attention, and Transformers outperform Smart Persistence, HC and NWP for point forecasts. LSTM, TCN, and TCN + Attention perform very well for point prediction. NGBoost performs comparably, or better, but falls behind in probabilistic evaluation. The probabilistic prediction results in Table 2 show that TCN (and TCN + Attention) obtain superior results against all benchmarks

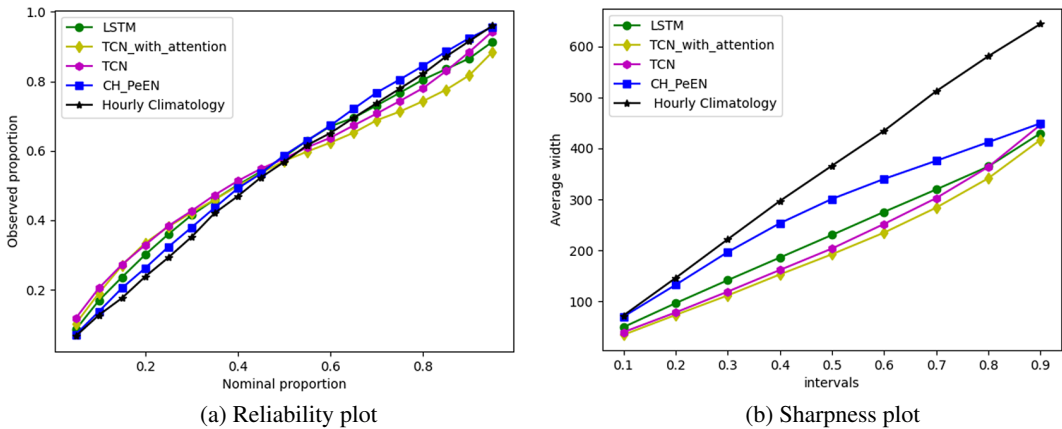


Figure 3. Reliability and sharpness plots at Penn State station.

except CH-PeEN in terms of the CRPS scores. Overall, LSTMs perform equally well. Transformers however do not come close to the other proposed models for both point and probabilistic evaluation (except for Desert Rock station). The CH-PeEN benchmark consistently performs slightly better than the best-performing probabilistic models. To investigate this, we refer to the reliability and sharpness diagrams for the station Penn State in Figure 3. We clearly note all our proposed models have better sharpness (as their curves are lower) in their forecasts than CH-PeEN, even though it very reliable.

6. Discussion

Our encouraging results demonstrate that deep sequence learning algorithms hold promise for producing improved week-ahead forecasts as they outperform most of the literature benchmarks. Our methods also provide a distribution over the prediction, and this additional knowledge of uncertainty can be extremely important in efficient power system and generator planning. Our proposed models outperform a machine learning-based approach (Duan et al., 2020) in probabilistic forecasting.

While temporal CNNs are a faster alternative to training LSTMs, they show an almost equal performance in this application, especially for probabilistic forecasts. The Attention mechanism proved useful when used in conjunction with TCNs but notably, not as much when we dispensed the convolution layers and used a transformer which is entirely based on attention.

Furthermore, as our part of our study, we wanted to look into the potential performance of our existing deep learning models if they are provided with an additional input feature of the NWP model ensemble. Tables 3 and 4 provide the results obtained when the 51 member ensemble is incorporated into our models. With the poor temporal resolution of these NWP predictors, we did not expect to see a huge performance improvement in the forecasts. We do observe an overall slight enhancement in performance with LSTM but not a clear trend with the TCN, TCN + Attention, and Transformer models. Further investigation is left to future work.

7. Conclusion

We provide a quantitative study and demonstrate the valuable potential of deep learning methods for 1 week-ahead solar irradiance forecasting, especially when such longer-term predictions are ill-served by existing NWP models. Week-ahead and longer forecasts, coupled with uncertainty estimates, can be very significant for future power systems operations when efficient energy planning will become increasingly important with greater shares of renewable energy penetrating into power systems. We hope this paper will encourage future work leveraging machine learning for long-term point and probabilistic forecasting, not only for solar power but also for other renewables and applications mitigating climate change.

Table 3. Results of the point forecasting pipeline with NWP ensemble included as features in our models.

	SP	HC	NGBoost	LSTM	TCN	TCN + Attention	Transformer
Sioux Falls, SD	222.34	220.14	171.53	187.37	183.71	189.97	215.57
Fort Peck, MT	203.11	225.12	156.82	147.31	164.02	153.48	176.53
Bondville, IL	269.6	230.99	196.36	193.85	212.27	198.57	229.75
Penn State, PA	235.57	225.67	186.48	196.78	206.61	213.77	218.41
Boulder, CO	242.01	220.52	177.53	183.75	196.69	190.98	224.79
Desert Rock, NV	146.83	196.15	111.83	104.68	120.23	130.51	125.34
Goodwin Creek, MS	252.01	230.06	180.41	190.9	209.99	206.94	218.45

Note. Results are in terms of RMSE scores.

Abbreviations: HC, hourly climatology; LSTM, long short-term memory; RMSE, root mean squared error; SP, smart persistence; TCN, temporal CNN.

Table 4. Results of the probabilistic forecasting pipeline with NWP ensemble included as features in our models.

	HC	CH-PeEN	NGBoost	LSTM	TCN	TCN + Attention	Transformer
Sioux Falls, SD	126.04	87.75	102.85	88.26	106.89	121.44	131.73
Fort Peck, MT	129.37	77.37	97.2	75.03	83.51	81.76	95.77
Bondville, IL	131.79	100.84	119.29	109.3	113.11	113.61	131.93
Penn State, PA	126.9	97.74	114.35	99.61	118.22	117.47	135.36
Boulder, CO	126.46	88.65	103.16	92.2	103.03	100.83	120.93
Desert Rock, NV	112.46	44.78	55.15	44.02	47.01	60.04	48.59
Goodwin Creek, MS	130.4	95.91	112.11	96.13	121.14	120.4	128.56

Note. Results are in terms of CRPS scores.

Abbreviations: CH-PeEN, complete history persistence ensemble; CRPS, continuous ranked probability score; HC, hourly climatology; LSTM, long short-term memory; NWP, numerical weather prediction; TCN, temporal CNN.

Author Contributions. Conceptualization: S.S., B.-M.H., C.M.; Data curation: S.S., B.-M.H.; Methodology: S.S., C.M.; Supervision: B.-M.H., C.M.; Writing—original draft: S.S., B.-M.H., C.M.; Writing—review and editing: S.S., B.-M.H., C.M. All authors approved the final submitted draft.

Competing Interests. C.M. is the Editor-in-Chief of Environmental Data Science. This paper was independently reviewed and accepted for publication.

Data Availability Statement. We use open-source NOAA's SURFRAD network (surface radiation budget network for atmospheric research, Augustine et al., 2000) that provides the ground-truth solar irradiance and meteorological measurements from seven sites across the US in different climatic zones (<https://gml.noaa.gov/grad/surfrad/>).

Funding Statement. This work received no specific grant from any funding agency, commercial or not-for-profit sectors.

Provenance. This article is part of the Climate Informatics 2022 proceedings and was accepted in *Environmental Data Science* on the basis of the Climate Informatics peer review process.

References

- Alzahrani A, Shamsi P, Dagli C and Ferdowsi M (2017) Solar irradiance forecasting using deep neural networks. *Procedia Computer Science* 114, 304–313.
- Augustine JA, DeLuisi JJ and Long CN (2000) Surfrad – A national surface radiation budget network for atmospheric research. *Bulletin of the American Meteorological Society* 81(10), 2341–2358.

- Bahdanau D, Cho K and Bengio Y** (2015) Neural machine translation by jointly learning to align and translate. In Bengio Y and LeCun Y (eds), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*. Available at <http://arxiv.org/abs/1409.0473> accessed on 28-01-2022.
- Borovykh A, Bohte S and Oosterlee CW** (2017) Conditional time series forecasting with convolutional neural networks. Preprint, arXiv:1703.04691.
- Brahma B and Wadhvani R** (2020) Solar irradiance forecasting based on deep learning methodologies and multi-site data. *Symmetry* 12(11), 1830.
- Doubleday K, Hernandez VVS and Hodge B-M** (2020) Benchmark probabilistic solar forecasts: Characteristics and recommendations. *Solar Energy* 206, 52–67.
- Duan T, Anand A, Ding DY, Thai KK, Basu S, Ng A and Schuler A** (2020) NGBoost: Natural gradient boosting for probabilistic prediction. In *International Conference on Machine Learning*. PMLR, pp. 2690–2700.
- Gensler A, Henze J, Sick B and Raabe N** (2016) Deep learning for solar power forecasting—An approach using autoencoder and lstm neural networks. In *IEEE International Conference on Systems, Man, And Cybernetics (SMC)*. IEEE, pp. 002858–002865.
- Gneiting T, Balabdaoui F and Raftery AE** (2007) Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69(2), 243–268.
- Granier C, Darras S, van der Gon HD, Jana D, Elguindi N, Bo G, Michael G, Marc G, Jalkanen J-P, Kuenen J, Lioussé C, Quack B, Simpson D, Sindelarova K** (2019) *The Copernicus Atmosphere Monitoring Service Global and Regional Emissions* (April 2019 version). PhD thesis, Copernicus Atmosphere Monitoring Service.
- Lauret P, David M and Pinson P** (2019) Verification of solar irradiance probabilistic forecasts. *Solar Energy* 194, 254–271.
- Lin Y, Koprinska I and Rana M** (2020) Temporal convolutional neural networks for solar power forecasting. In *International Joint Conference on Neural Networks (IJCNN)*. IEEE, pp. 1–8.
- Mishra S and Palanisamy P** (2018) Multi-time-horizon solar forecasting using recurrent neural network. In *2018 IEEE Energy Conversion Congress and Exposition (ECCE)*. IEEE, pp. 18–24.
- Qin Y, Song D, Chen H, Cheng W, Jiang G and Cottrell GW** (2017) A dual-stage attention-based recurrent neural network for time series prediction. In Sierra C (ed.), *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19–25*, pp. 2627–2633. <https://doi.org/10.24963/ijcai.2017/366>
- Saha M, Naik B and Monteleoni C** (2020) Probabilistic and point solar forecasting using attention based dilated convolutional neural network. In *EGU General Assembly Conference Abstracts*, p. 12818.
- Shaojie Bai JZK and Koltun V** (2018) An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. Preprint, arXiv:1803.01271.
- Song H, Rajan D, Thiagarajan JJ and Spanias A** (2018) Attend and diagnose: Clinical time series analysis using attention models. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Tuohy A, Zack J, Haupt SE, Sharp J, Ahlstrom M, Dise S, Gritm E, Mohrlen C, Lange M, Casado MG, Black J, Marquis M, Collier C.** (2015) Solar forecasting: Methods, challenges, and performance. *IEEE Power and Energy Magazine* 13(6), 50–59.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł and Polosukhin I** (2017) Attention is all you need. In Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S and Garnett R (eds), *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Available at <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf> accessed on 28-01-2022.
- Oord AVD, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior A and Kavukcuoglu K** (2016) Wavenet: A generative model for raw audio. Preprint, arXiv:1609.03499.
- Wang H, Lei Z, Zhang X, Zhou B and Peng J** (2019) A review of deep learning for renewable energy forecasting. *Energy Conversion and Management* 198, 111799.
- Wu N, Green B, Ben X and O'Banion S** (2020) Deep transformer models for time series forecasting: The influenza prevalence case. Preprint, arXiv:2001.08317.
- Zelikman E, Zhou S, Irvin J, Raterink C, Sheng H, Kelly J, Rajagopal R, Ng AY and Gagne D** (2020) Short-term solar irradiance forecasting using calibrated probabilistic models. Preprint, arXiv:2010.04715.
- Zhang H, Goodfellow I, Metaxas D and Odena A** (2019) Self-attention generative adversarial networks. In *International Conference on Machine Learning*. PMLR, pp. 7354–7363.