Adv. Stat. Clim. Meteorol. Oceanogr., 11, 89–105, 2025 https://doi.org/10.5194/ascmo-11-89-2025 © Author(s) 2025. This work is distributed under the Creative Commons Attribution 4.0 License.





Machine-learning-based probabilistic forecasting of solar irradiance in Chile

Sándor Baran¹, Julio C. Marín^{2,3}, Omar Cuevas^{3,4}, Mailiu Díaz⁵, Marianna Szabó¹, Orietta Nicolis⁵, and Mária Lakatos¹

¹Faculty of Informatics, University of Debrecen, Debrecen, Hungary
²Department of Meteorology, University of Valparaíso, Valparaíso, Chile
³Center for Atmospheric Studies and Climate Change (CEACC), University of Valparaíso, Valparaíso, Valparaíso, Chile
⁴Institute of Physics and Astronomy, University of Valparaíso, Valparaíso, Chile
⁵Faculty of Engineering, Andrés Bello University, Viña del Mar, Chile

Correspondence: Sándor Baran (baran.sandor@inf.unideb.hu)

Received: 5 December 2024 - Revised: 1 March 2025 - Accepted: 22 March 2025 - Published: 11 June 2025

Abstract. By the end of 2023, renewable sources covered 63.4 % of the total electric-power demand of Chile, and, in line with the global trend, photovoltaic (PV) power showed the most dynamic increase. Although Chile's Atacama Desert is considered to be the sunniest place on Earth, PV power production, even in this area, can be highly volatile. Successful integration of PV energy into the country's power grid requires accurate short-term PV power forecasts, which can be obtained from predictions of solar irradiance and related weather quantities. Nowadays, in weather forecasting, the state-of-the-art approach is the use of ensemble forecasts based on multiple runs of numerical weather prediction models. However, ensemble forecasts still tend to be uncalibrated or biased, thus requiring some form of post-processing. The present work investigates probabilistic forecasts of solar irradiance for regions III and IV in Chile. For this reason, eight-member short-term ensemble forecasts of solar irradiance for the calendar year 2021 are generated using the Weather Research and Forecasting (WRF) model; these are then calibrated using the benchmark ensemble model output statistics (EMOS) method based on a censored Gaussian law and its machine-learning-based distributional regression network (DRN) counterpart. Furthermore, we also propose a neural-network-based post-processing method, resulting in improved eightmember ensemble predictions. All forecasts are evaluated against station observations for 30 locations in the study area, and the skill of post-processed predictions is compared to the raw WRF ensemble. Our case study confirms that all studied post-processing methods substantially improve both the calibration of probabilistic forecasts and the accuracy of point forecasts. Among the methods tested, the corrected ensemble exhibits the best overall performance. Additionally, the DRN model generally outperforms the corresponding EMOS approach.

1 Introduction

According to the latest report of the International Renewable Energy Agency (IRENA, 2024), the largest ever increase in renewable-power capacity was observed in 2023, nearly 75 % of which was newly installed solar energy. As a result, by the end of 2023, the renewable-energy share had reached 43 % of the global installed power capacity, and this ratio was even higher in South America (71.4 %). In particular, renewable sources covered 63.4 % of the total electric-power de-

mand of Chile, 39.7% of which came from photovoltaic (PV) energy. In line with the global trend, with the addition of 1949 MW, in 2023, PV power accounted for the most substantial increase of 30.4%.

Although Chile's Atacama Desert is considered to be the sunniest place on Earth, with the highest long-term solar irradiance (Rondanelli et al., 2015), PV power production can be highly volatile, which raises a strong demand for accurate PV power forecasts from power grid operators. A standard approach to PV power forecasting is to consider global hor-

izontal irradiance (GHI) forecasts (and possibly forecasts of other weather variables) and to convert them into PV power with the help of a model chain; see, for example, Mayer and Yang (2022) or Horat et al. (2025). In the present study, we concentrate on solar irradiance predictions as they are highly correlated with the PV model chain outputs.

Traditionally, solar irradiance forecasts are obtained as outputs of numerical weather prediction (NWP) models, which describe the behaviour of the atmosphere with the help of partial differential equations. The state-of-the-art approach is to run these models simultaneously with various initial conditions and/or parameterizations, resulting in a probabilistic prediction as part of an ensemble forecast (Bauer et al., 2015; Buizza, 2018a). Nowadays, all major weather centres operate their own ensemble prediction systems (EPSs); one of the most prominent is the Integrated Forecast System (IFS) of the European Centre for Medium-Range Weather Forecasts (ECMWF), providing 51-member medium-range ensemble forecasts at 9 km resolution and 101-member extended-range forecasts at 32 km resolution (ECMWF, 2024). Nevertheless, in the last few years, NWP models gained strong competitors in the form of machinelearning-based, fully data-driven forecasts such as Pangu-Weather (Bi et al., 2023) or the more recent ECMWF Artificial Intelligence Forecasting System (AIFS; Lang et al., 2024), which was the first to issue AI-based ensemble predictions, followed by GenCast of Google DeepMind (Price et al., 2025).

Despite the efforts devoted to improving the EPSs, ensemble forecasts might exhibit deficiencies such as bias or lack of calibration, thus requiring some form of statistical postprocessing (Buizza, 2018b). In the last decades, many postprocessing methods have been suggested for a wide spectrum of weather variables; for an overview, see, for example, Vannitsem et al. (2021) or Schulz and Lerch (2022). Among these parametric approaches are the ensemble model output statistics method (EMOS; Gneiting et al., 2005) and the distributional regression network approach (DRN; Rasp and Lerch, 2018), which provide full predictive distributions in the form of a single parametric law. In the EMOS method, the parameters of the predictive distribution depend on the forecast ensemble via appropriate link functions, whereas, in the DRN approach, one trains a neural network, which connects the ensemble forecasts and possible other covariates to the distributional parameters. EMOS and DRN models for different weather quantities usually differ in terms of the parametric family describing the predictive distribution, and the EMOS link functions and the architectures of the DRN networks might also vary. Nonparametric methods include quantile regression, providing the predictive distribution in terms of its quantiles using either statistical tools (Friederichs and Hense, 2007; Bremnes, 2019) or machine learning techniques (Taillardat et al., 2016; Bremnes, 2020), and methods that directly improve the raw ensemble predictions, such as quantile mapping (Hamill and Scheuerer, 2018) or memberby-member post-processing (Van Schaeybroeck and Vannitsem, 2015).

To calibrate solar irradiance ensemble forecasts, Schulz et al. (2021) suggest EMOS models based on censored logistic and censored normal (CN0) distributions, while Baran and Baran (2024) and Horat et al. (2025) propose DRN counterparts of the latter. Furthermore, La Salle et al. (2020) compare linear quantile regression and the analogue ensemble technique to EMOS models utilizing truncated normal and truncated generalized extreme-value distributions. Bakker et al. (2019) consider quantile regression, quantile regression neural networks, random forests, and gradient-boosting decision trees, while Song et al. (2024) propose a non-crossing quantile regression neural network. According to the classification of Yang and van der Meer (2021), all the postprocessing methods mentioned above can be considered to be probabilistic-to-probabilistic (P2P) approaches; see also Yang and Kleissl (2024, Sect. 8.6).

The present work investigates probabilistic forecasts of solar irradiance for regions III and IV in Chile, covering an area with the second-largest PV power potential after the Atacama Desert (Molina et al., 2017). For this reason, eightmember ensemble forecasts of solar irradiance for the calendar year 2021 are generated using the Weather Research and Forecasting (WRF) model (Skamarock et al., 2019); these are then calibrated using the benchmark CN0 EMOS model and its DRN counterpart. Furthermore, we also propose a neural-network-based post-processing method, resulting in improved eight-member ensemble predictions. All forecasts are evaluated against station observations for 30 locations in the study area, and, in a detailed case study, the skill of postprocessed predictions is compared to the raw WRF ensemble. To the best of the authors' knowledge, no studies have been published yet that address the post-processing of solar irradiance ensemble forecasts for this part of the world. Thus, the main contributions of this work can be summarized in the following way.

- This paper provides a detailed analysis of the forecast skill of WRF probabilistic predictions of solar irradiance for two Chilean regions that are particularly important from the point of view of PV power production.
- This paper investigates the skill of state-of-the-art parametric post-processing methods.
- We also develop a machine-learning-based method for generating improved irradiance ensemble forecasts and assess its predictive performance.

The paper is organized as follows. The description of the WRF model configurations and the observed data and a preliminary assessment of the performance of WRF forecasts are given in Sect. 2. Section 3 introduces the applied postprocessing methods, the approaches to training-data selection for modelling, and the considered forecast evaluation tools. The results of our case study are reported in Sect. 4, and the paper concludes with a summary and discussion in Sect. 5.

2 Data

2.1 WRF model configuration and ensemble members

The Advanced Research module of the Weather Research and Forecasting (WRF) model, version 4.4.2 (Skamarock et al., 2019), was used in this study. We consider eightmember ensemble forecasts for solar irradiance (downward shortwave flux given in $W m^{-2}$) for the calendar year 2021 as generated by the model. All forecasts are initialized at 00:00 UTC with a 1 h temporal resolution and a forecast horizon of 48 h. The eight simulations were configured with two nested domains, as depicted in Fig. 1a, with 9 km (d01) and 3 km (d02) horizontal resolutions, 36 vertical levels at variable resolutions, and enhanced density near the surface and the tropopause. We will show only the results for the smallest domain (d02) with the highest resolution (3 km), provided separately in Fig. 1b. This domain includes regions III and IV of Chile, covering several coastal and interior towns and cities to the west and part of the Andes Cordillera to the east. Forecasts from the Global Forecast System (GFS) model at $0.25^{\circ} \times 0.25^{\circ}$ horizontal resolution provided the initial and boundary conditions for the WRF model regional forecasts every 3 h. The WRF simulations employed land use data based on the Moderate Resolution Imaging Spectroradiometer (MODIS) at 15 arcsec (approximately 0.4 km in horizontal resolution). We also set a two-way interaction, which allows the inner domain to provide feedback to its parent domain.

The eight simulations used the Noah-MP land surface model (Niu et al., 2011; Yang et al., 2011) to parameterize the surface-atmosphere interaction in all domains. The Noah-MP scheme forecasts the soil temperature and moisture and provides fractional snow cover and frozen-soil physics. Convective processes in all domains were calculated with the Kain–Fritsch scheme (Kain–F; Kain, 2004), while the Thompson double-moment scheme (Thompson et al., 2008) was used for microphysics. The eight-member simulations differ in terms of the employed radiation and planetary boundary layer (PBL) scheme. Four radiation and five PBL schemes were combined to form the eight-member simulations, whose descriptions are displayed in Table 1. The Yonsei University model (YSU; Hong et al., 2006), the Mellor-Yamada-Janjić model (MYJ; Janjić, 1994), the quasi-normal-scale elimination model (QNSE; Sukoriansky et al., 2005), the University of Washington model (UW; Bretherton and Park, 2009), and the asymmetric convective model (ACM2; Pleim, 2007) were used to calculate the PBL processes, whereas the rapid radiative transfer model (RRTMG; Iacono et al., 2005), the Dudhia model (Dudhia, 1989), the Fu-Liou-Gu model (FLG; Gu et al., 2011), and **Table 1.** Description of radiation (RAD), planetary boundary layer (PBL), land surface model (LSM), cumulus, and microphysics (MP) parameterizations used for each of the eight WRF ensemble members.

Memb.	PBL	RAD	LSM	Cumulus	MP
1	YSU	RRTMG	Noah-MP	Kain–F	Thompson
2	YSU	Dudhia	Noah-MP	Kain–F	Thompson
3	YSU	FLG	Noah-MP	Kain–F	Thompson
4	YSU	Goddard	Noah-MP	Kain–F	Thompson
5	MYJ	RRTMG	Noah-MP	Kain–F	Thompson
6	QNSE	RRTMG	Noah-MP	Kain–F	Thompson
7	UW	RRTMG	Noah-MP	Kain–F	Thompson
8	ACM2	RRTMG	Noah-MP	Kain–F	Thompson

the new Goddard scheme (Chou et al., 1999, 2001) were used to parameterize the longwave and shortwave radiation processes. Note that one could have considered more combinations of PBL and radiation schemes, resulting in more ensemble members. However, on the one hand, we had limited computational resources, which posed constraints on both the ensemble size and the horizontal resolution. On the other hand, the current study is part of a broader research project dealing with the calibration of renewable-energyrelated quantities (solar irradiance, wind speed, precipitation accumulation) in Chile, and we wanted to ensure comparability with our earlier results on post-processing eightmember WRF temperature (Díaz et al., 2020) and wind speed (Díaz et al., 2021) predictions. The relatively small ensemble size might affect the variability of the raw WRF forecasts; nonetheless, it does not affect the choice of the possible postprocessing methods.

2.2 Solar irradiance observations

Solar irradiance observations (given in W m⁻²) are provided by the Chilean national weather service (DMC, an abbreviation for the Spanish name). We consider data from 30 stations in the Atacama and Coquimbo regions, between the coast and the Andes Cordillera. Table 2 describes each station's name and location (latitude, longitude, altitude), and Fig. 1b shows their spatial distribution over the study region. The observations with hourly temporal resolution were downloaded from the website: https://climatologia.meteochile.gob.cl/ (last access: 1 March 2025).

2.3 Forecast skill of the WRF ensemble

The matching of the WRF ensemble predictions with observations at the monitoring stations listed in Table 2 is performed by extracting forecasts for the nearest grid points from the WRF domain d02. As depicted in Fig. 1b, the topography of this region is rather complex; the station altitudes range from 80 to 2154 m, and the differences in terms of station elevation have an impact on the forecast perfor-



Figure 1. (a) The two WRF nested-domain configurations (d01 and d02) used for each simulation and the region's topography (shaded colours). (b) Zoomed-in image of domain d02, showing the location of the 30 observation stations (red circles).

Table 2. The name, longitude, latitude, altitude, and region of 30 irradiance observation stations used in this study.

No.	Name	Longitude	Latitude	Altitude (m)	Region
1	Desierto de Atacama, Caldera Ad.	-70.781	-27.254	197	III
2	Universidad de Atacama	-70.353	-27.359	362	III
3	Amolana	-70.010	-27.960	1090	III
4	Copiapó	-70.408	-27.35	574	III
5	Tierra Amarilla, T. Lautaro	-70.000	-28.976	1173	III
6	Tierra Amarilla, Jotabeche	-70.245	-27.589	601	III
7	La Copa	-70.622	-27.346	188	III
8	Freirina, Nicolasa	-71.01	-28.517	156	III
9	CE Huasco	-70.798	-28.581	470	III
10	Freirina, Vallenar	-70.943	-28.526	230	III
11	Alto del Carmen	-70.449	-28.768	822	III
12	Freirina	-71.104	-28.507	100	III
13	La Florida, La Serena Ad.	-71.207	-29.914	137	IV
14	La Higuera, El Trapiche	-71.116	-29.372	281	IV
15	Ovalle Escuela Agrícola	-71.187	-30.580	310	IV
16	Punitaqui (FDF)	-71.256	-30.780	216	IV
17	El Tololo	-70.804	-30.168	2154	IV
18	Algarrobo Bajo	-71.451	-30.633	80	IV
19	Camarico	-71.322	-30.699	290	IV
20	El Palqui	-70.927	-30.774	504	IV
21	Vicuña, Los Pimientos	-70.697	-30.034	642	IV
22	Monte Patria, Municipalidad	-70.958	-30.698	406	IV
23	Paiguano	-70.511	-30.047	1222	IV
24	Ovalle Recoleta	-71.166	-30.491	403	IV
25	Monte Patria	-70.939	-30.682	563	IV
26	Chaguaral	-70.731	-30.857	1194	IV
27	La Polvareda	-71.240	-30.883	265	IV
28	Ajial de Quiles	-71.384	-30.920	464	IV
29	Canela Baja (Liceo P. José Herde)	-71.413	-31.390	343	IV
30	Liceo Samuel Román Rojas	-70.999	-31.188	936	IV

mance. Figure 2 shows the boxplots of the median forecast error of the WRF ensemble (median deviation of the eight individual ensemble members from the verifying observation, that is, the median bias of the WRF members) at the various observation stations for the whole calendar year 2021 for lead times corresponding to the period between 12:00 and 00:00 UTC, when positive irradiance is likely to be observed. Raw WRF forecasts systematically overestimate the actual irradiance, and both the magnitude of the bias and the spread of the median forecast error strongly depend on the loca-



Figure 2. Median forecast error $(W m^{-2})$ of the WRF ensemble at the observation stations listed in Table 2 for all dates and lead times of 12–24 h and 36–48 h.



Figure 3. Forecast error $(W m^{-2})$ of the individual ensemble members for all dates and locations for lead times of (a) 12–24 h and (b) 36–48 h.



Figure 4. Forecast error $(W m^{-2})$ of the ensemble median for all dates and locations as functions of the lead time.

tion. However, in contrast to Díaz et al. (2020), where the station altitude played a key role in the capability of WRF temperature forecasts, here, despite the WRF irradiance ensemble performing the best at the highest station (no. 17; El Tololo), one cannot find a clear connection between the elevation and the forecast skill. The same positive bias can be observed in Fig. 3, displaying the boxplots of the fore-

cast error of the individual ensemble members for all dates and locations, treating shorter (12–24 h) and longer (36–48 h) lead times separately. Although longer forecast horizons result in slightly larger forecast errors, Fig. 3a and b convey the same message. Ensemble member 2 substantially outperforms the other seven members by exhibiting the smallest bias; the largest error corresponds to member 3, followed by member 4, whereas the performance of members 1 and 5– 6 is fairly similar. Finally, the boxplots of the diurnal evolution of the forecast error of the ensemble median depicted in Fig. 4 indicate that the positive bias is also systematic along the forecast horizons of 12–24 h and 36–48 h, when positive irradiance is likely to be observed, and the largest errors correspond to 15:00 and 16:00 UTC, when the solar irradiance reaches its peak.

3 Post-processing and forecast evaluation methods

As mentioned in the Introduction, we consider two different parametric methods for post-processing WRF solar irradiance ensemble forecasts. The key step in parametric modelling is the choice of the predictive distribution. Addressing non-negativity of solar irradiance, several parametric models utilize distributions left truncated from below at zero; see, for example, La Salle et al. (2020) or Yang (2020). However, such models require the specification of times of day with positive irradiance, whose periods depend strongly on the location and season. This deficiency can be solved by considering laws that assign a positive mass to the event of zero irradiance as such a distribution can handle even the night hours, when both predicted and observed irradiance are zero. A popular choice is to left censor a suitable distribution at zero, which proved to be successful a successful approach, e.g. in modelling precipitation accumulation (Scheuerer, 2014; Baran and Nemoda, 2016). Here, we utilize the censored normal EMOS approach of Schulz et al. (2021), considered to be a benchmark method, and the corresponding DRN model investigated by Baran and Baran (2024) and Horat et al. (2025). Parametric models based on this particular predictive distribution demonstrated excellent skill for various ensemble forecasts covering different geographical areas.

Beyond parametric modelling, we also suggest a neuralnetwork-based distribution-free ensemble correction technique, where the output is a calibrated forecast ensemble. Although, in this post-processing approach, there is no restriction on the number of generated corrected forecasts, for the sake of fair comparability, we consider the ensemble size of eight of the raw WRF ensembles.

Compared to the EMOS and DRN models, the aforementioned method lacks the advantage of providing a full predictive distribution; however, similarly to the nonparametric post-processing techniques, it allows more flexible modelling, and the output can be interpreted in the same way as the raw ensemble forecast.

In what follows, let f_1, f_2, \ldots, f_8 denote the eightmember WRF irradiance forecast of a given forecast horizon for a given location and time, and let \overline{f} and S^2 denote the corresponding ensemble mean and variance, respectively.

3.1 EMOS model for solar irradiance

Let $G(x|\mu, \sigma)$ denote the cumulative distribution function (CDF) of a Gaussian distribution with mean μ and standard deviation $\sigma > 0$:

$$G(x|\mu,\sigma) := \Phi\left(\frac{x-\mu}{\sigma}\right), \ x \in \mathbb{R},$$

with Φ denoting the CDF of a standard normal law. Then the CDF of a normal distribution with location μ and scale σ left censored at zero (CN0) equals

$$G_0^c(x|\mu,\sigma) := \begin{cases} G(x|\mu,\sigma), & x \ge 0, \\ 0, & x < 0 \end{cases}$$
(1)

This distribution assigns mass $G_0^c(0|\mu, \sigma)$ to the origin and has a mean of

$$\kappa = \mu \Phi(\mu/\sigma) + \sigma \varphi(\mu/\sigma),$$

where φ denotes the standard normal probability density function (PDF). Following Schulz et al. (2021) and Baran and Baran (2024), the ensemble members are connected to the parameters of the CN0 distribution via equations (link functions):

$$\mu = \gamma_0 + \gamma_1 f + \gamma_2 p_0$$
 and $\sigma = \exp(\delta_0 + \delta_1 \log S)$,

where p_0 is the proportion of WRF ensemble members predicting zero irradiance; that is,

$$p_0 := \frac{1}{8} \sum_{k=1}^{8} \mathbb{I}_{\{f_k=0\}}$$

with \mathbb{I}_H denoting the indicator function of a set H. Model parameters γ_0 , γ_1 , γ_2 , δ_0 , and $\delta_1 \in \mathbb{R}$ are estimated following the optimum score principle of Gneiting and Raftery (2007) that is obtained by minimizing the mean value of a proper scoring rule (in our case, the continuous ranked probability score (CRPS) defined by Eq. (2) in Sect. 3.5) over appropriate training data comprising past forecast–observation pairs. This general approach is the most standard in parametric ensemble post-processing, where replacing the CRPS with the ignorance score, defined as the negative logarithm of the predictive PDF evaluated for the corresponding observation (see, for example, Wilks, 2019, Sect. 9.5.3), results in the conventional maximum-likelihood estimates.

3.2 DRN model for solar irradiance

Distributional regression networks (DRNs), first employed by Rasp and Lerch (2018) for calibrating temperature ensemble forecasts, represent an advanced class of machinelearning-based post-processing models. Similarly to the EMOS approach, they extend traditional regression techniques by predicting parameters of the forecast distribution belonging to a given family, in our case, the location μ and scale σ of the CN0 distribution specified by the CDF (Eq. 1). The DRN approach enhances the statistical calibration of ensemble forecasts, providing a more comprehensive understanding of prediction uncertainty. DRNs typically leverage predictor variables, such as NWP quantities and station characteristics, to inform their predictions. Additionally, station embeddings enable the network to recognize location-specific information, capturing unique features and patterns relevant to individual stations.

DRNs are often implemented using a multilayer perceptron (MLP; Goodfellow et al., 2016) architecture, which improves their ability to capture complex relationships between the NWP forecasts and other covariates and the distributional parameters to be predicted. The feed-forward structure of the network allows input features to pass through multiple hidden layers, where neurons apply weighted transformations and activation functions to the input signals. Activation functions such as ReLU (rectified linear unit), sigmoid, or tanh introduce non-linearity, enabling the model to learn intricate patterns and relationships in the data. Similarly to EMOS modelling, the weights of the MLP neural networks connecting the input covariates with the distributional parameters μ and σ are estimated by minimizing the mean CRPS of the CN0 predictive distributions over the training data, optimizing probabilistic forecasts to align with observed outcomes. Note that, building upon the work of Rasp and Lerch (2018), which focused on fully connected networks such as MLP, recently, convolutional layers have also gained popularity within DRNs (Veldkamp et al., 2021; Li et al., 2022). These layers leverage local patterns and spatial hierarchies in the data, further enhancing the model's ability to capture complex relationships. However, in this study, we keep the traditional MLP architecture, consisting of an input layer with as many neurons as the number of input covariates, a few hidden layers (shallow network), and an output layer with two neurons corresponding to the distributional parameters. Further details regarding the applied MLP network can be found in Sect. 3.6.

To further enhance model performance and to prevent overfitting, early stopping is often employed as a regularization technique during training. By monitoring performance based on a validation dataset, training halts when improvements cease, ensuring the model retains its generalization capabilities. This approach not only optimizes the training process but also contributes to more robust predictions.

3.3 Machine-learning-based forecast improvement

In addition to constructing predictive distributions, another possible alternative is statistical post-processing using machine learning models that generate improved ensemble forecasts. For this purpose, similarly to the DRN model described in Sect. 3.2, we use a neural network based on the MLP architecture; however, here, the number of neurons in the output layer equals the number of ensemble members to be generated. The other principal difference is the implementation of the loss function. Due to the nature of the output, here, we apply the ensemble CRPS given by Eq. (3) (see Sect. 3.5), with the constraint that the predicted solar irradiance forecasts can only be non-negative. This post-processing method is flexible as it is distribution-free, and the number of ensemble members to be generated is up to the user, provided enough training data are available. However, to ensure direct comparability with the WRF forecasts, we create an eightmember prediction, referred to as the corrected ensemble, for each verification day, location, and forecast horizon.

3.4 Training-data selection

The efficiency of all post-processing methods, including the ones described in Sect. 3.1–3.3, strongly depends on the spatial and temporal decomposition of the training data.

From the point of view of temporal selection, a popular approach is using a sliding window, where the model for a given date is trained using forecast–observation pairs from the preceding *n* calendar days. This simple method, also utilized in our study, allows for a quick adaptation to seasonal variations or model changes; nevertheless, larger time shifts, such as monthly, seasonal, or yearly windows, might also be beneficial (see, for example, Jobst et al., 2023). Alternatively, one can consider a fixed, very long training period, a popular approach in machine-learning-based post-processing methods requiring a large amount of training data (see, for example, Schulz and Lerch, 2022; Horat et al., 2025). For a systematic comparison of time-adaptive model training schemes, we refer the reader to Lang et al. (2020).

Regarding the spatial composition of training data, the traditional approaches are local and regional (global) modelling (Thorarinsdottir and Gneiting, 2010). Local models are based on past forecast-observation pairs for the actual location under consideration, whereas the regional approach utilizes historical data of all studied stations. In general, local models outperform their regional counterparts, provided one has enough location-specific training data to avoid numerical issues during the calibration process. The regional approach, resulting in a single set of EMOS parameters or neural network weights for the whole ensemble domain, is more suitable if only short training periods are allowed; however, it might hardly handle large and heterogeneous areas. To combine the advantageous properties of the above two approaches to spatial selection, Lerch and Baran (2017) suggested a novel clustering-based semi-local method which appeared to be successful for several different weather variables and ensemble domains; see, for example, Baran et al. (2020), Szabó et al. (2023), or Baran and Lakatos (2024). For a given date of the verification period, a feature vector is assigned to each observation station, representing the station climatology and the forecast error of the ensemble mean during the training period. Based on these feature vectors, the stations are then grouped into clusters using k-means clustering (see, for example, Wilks, 2019, Sect. 16.3.1), and, within each cluster, regional modelling is performed. In the case of rolling training windows, the stations are regrouped dynamically for each particular training set.

3.5 Forecast verification

As mentioned, the parameters of the CN0 EMOS model described in Sect. 3.1 are estimated by minimizing the mean of a proper scoring rule over the training data. In particular, we consider the continuously ranked probability score (CRPS; Wilks, 2019, Sect. 9.5.1), which is probably the most popular proper verification score in atmospheric sciences. Given a probabilistic forecast expressed as a predictive CDF *F* and an observation $x \in \mathbb{R}$, the CRPS is defined as

$$CRPS(F, x) := \int_{-\infty}^{\infty} \left[F(y) - \mathbb{I}_{\{y \ge x\}} \right]^2 dy$$
$$= E|X - x| - \frac{1}{2}E|X - X'|, \qquad (2)$$

where X and X' are independent random variables distributed according to F and with a finite first moment. The CRPS is a negatively oriented score (smaller values indicate better performance), and it simultaneously assesses both the calibration and the sharpness of a probabilistic forecast. Calibration means a statistical consistency between forecasts and observations, while sharpness refers to the concentration of the probabilistic forecasts. Furthermore, the expression of the CRPS on the right-hand side of Eq. (2) indicates that it can be expressed in the same unit as the observation. For the censored normal distribution considered in Sect. 3.1 and 3.2, the CRPS has a closed form (see, for example, Jordan et al., 2019), which allows for a computationally efficient numerical optimization, hence making it eligible to serve as a loss function in both EMOS and DRN modelling. For a forecast ensemble f_1, f_2, \ldots, f_K , one should consider the empirical CDF \hat{F}_K , resulting in the expression

$$\operatorname{CRPS}(\hat{F}_K, x) = \frac{1}{K} \sum_{k=1}^{K} \left| f_k - x \right| - \frac{1}{2K^2} \sum_{k=1}^{K} \sum_{\ell=1}^{K} \left| f_k - f_\ell \right|.$$
(3)

See, for example, Krüger et al. (2021). The same definition applies to the corrected ensemble forecast produced by the approach described in Sect. 3.3 and to the calibrated samples generated from the EMOS or DRN predictive distributions. Note that the above formula differs slightly from the ensemble CRPS given in Wilks (2019, Sect. 9.7.3); however, in the R package scoringRules (Jordan et al., 2019), the expression of Eq. (3) is implemented.

In the case study of Sect. 4, the predictive performance of the competing probabilistic forecasts with a given forecast horizon is compared with the help of the mean CRPS over all forecast cases in the verification period. Furthermore, we also quantify the improvement in the mean CRPS of a probabilistic forecast F for a reference forecast F_{ref} using the continuous ranked probability skill score (CRPSS; see, for example, Gneiting and Raftery, 2007). This positively oriented quantity (the larger, the better) is defined as

$$CRPSS := 1 - \frac{CRPS_F}{\overline{CRPS}_{F_{ref}}},$$

where $\overline{\text{CRPS}}_F$ and $\overline{\text{CRPS}}_{F_{\text{ref}}}$ denote the mean CRPS corresponding to forecasts *F* and *F*_{ref}, respectively. In particular, in Sect. 4, for the reference prediction *F*_{ref}, we consider the raw WRF ensemble, while forecast *F* can be any of the postprocessed predictions described in Sect. 3.1–3.3.

A separate assessment of the calibration and sharpness of predictive distributions can be obtained with the help of the coverage and average width of $(1 - \alpha)100$ %, $\alpha \in (0, 1)$, central prediction intervals, respectively. By coverage, we mean the proportion of observations located between the $\alpha/2$ and $1 - \alpha/2$ quantiles of the predictive CDF, where properly calibrated forecasts result in values of around $(1 - \alpha)100\%$ (see, for example, Gneiting and Raftery, 2007). For a k-member forecast ensemble, one usually considers the nominal coverage (K - 1)/(K + 1)100% of the ensemble range (77.78%) for the eight-member WRF ensemble at hand), which is the probability of the rank of the observation for a calibrated prediction being greater than 1 and less than K + 1. Choosing α to match this nominal coverage allows for a fair comparison of ensemble forecasts with forecasts provided as full predictive distributions.

A further plausible tool for evaluating the calibration of ensemble forecasts is the verification rank histogram (or Talagrand diagram). The Talagrand diagram displays the ranks of the verifying observation with respect to the corresponding ensemble prediction (Wilks, 2019, Sect. 9.7.1), which, for a calibrated *k*-member ensemble, should follow a discrete uniform law on the set $\{1, 2, ..., K + 1\}$. The shape of a rank histogram reflects the source of the lack of calibration: \cup - and \cap -shaped histograms refer to underdispersion and overdispersion, respectively, whereas biased forecasts result in triangular shapes. Furthermore, the deviation of the distribution of the verification ranks from the uniform law can be quantified with the help of the reliability index:

$$\mathrm{RI} := \sum_{r=1}^{K+1} \Big| \rho_r - \frac{1}{K+1} \Big|,$$

where ρ_r denotes the relative frequency of rank *r* over all forecast cases in the verification period (Delle Monache et al., 2006). The continuous counterpart of the Talagrand diagram is the probability integral transform (PIT; Wilks, 2019, Sect. 9.7.1) histogram. The PIT is defined as the value of predictive CDF for the verifying observation, with possible randomization at the points of discontinuity. For a calibrated

predictive distribution, PIT is standardly uniform, and the interpretation of the various deviations of the shapes of the PIT histograms from uniformity is similar to that for the Talagrand diagrams.

Furthermore, the accuracy of point forecasts such as the forecast median is evaluated with the mean absolute error (MAE) as the median minimizes this score (Gneiting, 2011).

Finally, the statistical significance in score differences is assessed by accompanying some skill scores and score differences by 95 % block bootstrap confidence intervals. We consider 2000 samples calculated using the stationary bootstrap scheme with block lengths following a geometric distribution with a mean proportional to the cube root of the length of the investigated time interval (Politis and Romano, 1994).

3.6 Modelling and implementation details

In the case of the EMOS modelling, all 48 forecast horizons are treated separately, and model parameters are estimated using a clustering-based semi-local approach. Similarly to Lerch and Baran (2017), we consider 24-D vectors, where half of the features comprise equidistant quantiles of the climatological CDF over the training period and the other half consist of equidistant quantiles of the empirical CDF of the forecast error of the ensemble mean. After testing several combinations of training-period length and number of clusters, an 85 d rolling window is chosen, and, in general, six clusters are formed (an average of 425 forecast–observation pairs for each estimation task), provided each cluster contains at least three observation stations. Otherwise, the number of clusters is reduced, which might result in regional modelling.

Machine-learning-based forecasts are estimated regionally, and a single MLP neural network is trained for all lead times. CN0 DRN models use 20 d rolling training windows, while corrected forecasts depend on forecast–observation pairs of the preceding 25 d. These training-period lengths are the results of detailed data analysis.

Through testing various hidden-layer configurations for the CN0 DRN approach, we identified an optimal MLP model with two hidden layers, each containing 255 neurons, and considered a batch size of 1200. To enhance numerical stability, we standardize input features and remove missing data. The model is trained with the "Adam" optimizer, a learning rate of 0.01, and the ReLU activation function. Input features include the ensemble mean \overline{f} and variance S^2 , the proportion p_0 of zero-irradiance forecasts, station coordinates (latitude, longitude, altitude), and lead times. Feature importance was also assessed to evaluate each input's impact on model performance. The model output directly represents the parameters μ and σ of the CN0 predictive distribution, with only a squaring being applied to the scale parameter to ensure non-negativity. The training is capped at 500 epochs, but an early stopping callback – using 20 % of the training data as a validation set – often enabled convergence at approximately 50 epochs for the first verification day, with fewer epochs being required on subsequent days. To address the randomness inherent in training, for each forecast case, we train the network 10 times, deriving final predictions by averaging the distribution parameters across these sessions.

Our detailed tests suggest that, when DRN input features are restricted to ensemble statistics used in EMOS modeling (i.e. the same set of explanatory variables), the DRN's performance generally aligns with that of EMOS. This indicates that, without additional input information beyond the functionals of the ensemble forecasts, DRNs may have limited capacity for performance gains over EMOS in terms of accuracy and calibration. Notably, including station location information contributed significantly to enhanced model performance, which aligns with the findings of Horat et al. (2025).

For the MLP network providing corrected forecasts, we also comprehensively tested different input features and hyperparameters, though optimal results were achieved with the same settings as for the CN0 DRN approach. Due to the increase in the neurons in the output layer from two to eight, the number of weights to be estimated increased from 67 832 to 69368, which explains the 5 d increase in the trainingperiod length. Furthermore, to ensure the non-negativity of the generated forecasts, we take the maximum of zero and the predicted value, which is consistent with the handling of negative predictions in the loss function. Finally, as before, to account for the stochastic nature of the model, 10 independent runs are performed for each forecast case. The generated forecasts are then sorted in ascending order, and the average is computed over the sorted values. Note that, without sorting, the averaging would decrease the spread of the obtained eight-member forecast, which would no longer bear the same distributional characteristics as the originally generated eight-member corrected ensembles.

From the point of view of computational costs, which are negligible compared with the cost of generating the raw ensemble forecasts, there are no substantial differences between the three investigated post-processing approaches. The two machine-learning-based methods have similar architectures and rely on the same type of training data to estimate the weights of a single neural network for all locations and lead times. In general, the optimization step in EMOS modelling is much faster than in network training; however, separate semi-local modelling for all lead times usually requires 6×48 individual EMOS models.

4 Results

In the following case study, the predictive performances of the CN0 EMOS and CN0 DRN approaches described in



Figure 5. (a) Mean CRPS of raw and post-processed irradiance forecasts and (b) CRPSS of post-processed forecasts with respect to the raw WRF ensemble as functions of the lead time.

Sect. 3.1 and 3.2, respectively, are evaluated together with the machine-learning-based forecast improvement method presented in Sect. 3.3. To ensure a fair comparison of the latter with the EMOS and DRN models resulting in full predictive distributions, we also investigate the forecast skill of eight-member samples generated from the corresponding CN0 distributions, where one can consider their 1/9, $2/9, \ldots, 8/9$ quantiles, simple random sampling, or stratified samples (Hu et al., 2016). As preliminary analysis reveals only minor differences in the skill of equidistant quantiles and stratified samples, both outperforming random sampling, here, we work with equidistant quantiles of CN0 EMOS and CN0 DRN predictive distributions. The corresponding forecasts are referred to as CN0 EMOS-Q and CN0 DRN-Q, respectively.

In the following analysis, the predictive performance of the competing post-processing methods is tested on forecast– observation pairs for the 9 months of 1 April–31 December 2021 (275 calendar days, right after the first 85 d training window used in EMOS modelling), and the raw WRF irradiance forecasts are used as a reference.

Figure 5a displays the mean CRPS of raw and postprocessed WRF irradiance forecasts over all 30 locations and all 275 d in the verification period. Between 12:00 and 00:00 UTC (lead times of 12-24 h and 36-48 h), when positive irradiance is likely to be observed, all calibration methods outperform the raw ensemble by a wide margin. According to the CRPSS values of Fig. 5b, at this time of the day, the improvement with respect to the raw ensemble is around 50 % for all competing forecasts, and the ranking of the various methods is consistent for all forecast horizons. There are only minor differences in terms of skill between forecasts resulting in full predictive distributions (CN0 EMOS and CN0 DRN) and their sample counterparts (CN0 EMOS-Q and CN0 DRN-Q, respectively). The corrected ensemble results in the highest CRPSS, followed by the DRN and EMOS predictions. Note that, for 1-11 h and 25-35 h forecasts, the

Table 3. Overall mean CRPS of post-processed irradiance forecasts as a proportion of the mean CRPS of the raw WRF ensemble for observed irradiance of no less than 7.5 W m^{-2} .

CN0	CN0	CN0	CN0	Corr.
EMOS	EMOS-Q	DRN	DRN-Q	ensemble
52.58 %	53.69 %	47.90 %	49.04 %	44.57 %

skill score values are irrelevant; they are just results of numerical issues.

The clear ranking of the various calibrated forecasts is also confirmed by Table 3, summarizing their mean CRPS as a proportion of the mean CRPS of the WRF ensemble for forecast cases where the observed irradiance is at least $7.5 \,\mathrm{W}\,\mathrm{m}^{-2}$. Note that, from the point of view of PV energy production, only these cases are of any interest, and the threshold coincides with the one considered in Baran and Baran (2024) and suggested by the forecasters of the Hungarian Meteorological Service (HMS). One should also remark that the improvements in mean CRPS over the raw WRF ensemble provided in Table 3 are much higher than the gains due to CN0 EMOS and CN0 DRN post-processing of 11member AROME-EPS forecasts (Jávorné Radnóczi et al., 2020) of the HMS, as reported by Baran and Baran (2024, Table 3). An even smaller advantage of the post-processed forecasts over the raw 40-member ICON-EPS predictions (Zängl et al., 2015) of the German Meteorological Service (DWD, Deutsche Wetterdienst) was observed by Schulz et al. (2021).

To investigate the statistical significance of the differences in terms of the mean CRPS, in Fig. 6, the corresponding 95 % confidence intervals are added to the CRPSS values of the CN0 EMOS and CN0 DRN approaches. For lead times corresponding to the 12:00–00:00 UTC interval, neither the difference between the corrected ensemble and the CN0 DRN nor the deviation of the CN0 EMOS and the CN0 DRN is significant; however, between 13–20 h and 37–44 h, the best-



Figure 6. CRPSS with respect to the raw WRF ensemble of the CN0 EMOS and CN0 DRN models (together with 95 % confidence intervals) and the corrected ensemble as functions of the lead time.



Figure 7. (a) Coverage and (b) average width of the nominal 77.78 % central prediction intervals of post-processed and raw irradiance forecasts as functions of the lead time. In (a), the ideal coverage is indicated by the horizontal dotted line.

performing corrected ensemble significantly outperforms the worst-performing CN0 EMOS.

The improved calibration of post-processed forecasts can also be observed in Fig. 7a, displaying the coverage of the nominal 77.78 % central prediction intervals. To ensure comparability with the raw ensemble, for calibrated eightmember forecasts, we consider the coverage of the whole ensemble range. In the case of CN0 EMOS-Q and CN0 DRN-Q predictions based on equidistant quantiles of the corresponding predictive distributions, the ensemble ranges coincide with the 77.78 % central prediction intervals of the CN0 EMOS and CN0 DRN models, respectively, and so they are excluded from the analysis. At the hours of the highest solar irradiance (13:00-22:00 UTC), the coverage of the raw WRF forecasts is below 40%, where post-processed forecasts result in coverage values much closer to the targeted 77.78 %. A possible ranking of the three investigated methods can be formed with the help of Table 4, providing the mean absolute deviation in coverage from the nominal level for 13:00-22:00 UTC observations. According to Fig. 7b, the

Table 4. Mean absolute deviation in coverage from the nominal 77.78 % level over lead times corresponding to 13:00–22:00 UTC observations.

CN0 EMOS	CN0 DRN	Corr. ensemble	Ensemble
3.49 %	13.84 %	5.18 %	53.22 %

improved coverage of calibrated forecasts is accompanied by wider prediction intervals; however, the CN0 EMOS prediction resulting in the best coverage values is considerably sharper than the second-best corrected ensemble.

The verification rank and PIT histograms of Fig. 8, where we, again, consider only lead times corresponding to 12:00– 00:00 UTC observations, also confirm the positive effect of statistical post-processing. First, note that the PIT histograms of the CN0 predictive distributions and the rank histograms of the corresponding samples are almost identical, which agrees with the matching CRPS and CRPSS values of Fig. 5.



Figure 8. PIT histograms of CN0 EMOS and CN0 DRN predictive distributions and rank histograms of raw and post-processed eightmember irradiance forecasts, together with the corresponding reliability indices for lead times of 12–24 h and 36–48 h.

The highly U-shaped and asymmetric rank histograms of the raw WRF forecasts indicate strong underdispersion (the spread of the forecast is too low to represent the forecast error correctly) and a pronounced positive bias, respectively, resulting in reliability indices of 1.0210 (12-24 h forecasts) and 0.9779 (36-48 h forecasts). These deficiencies are substantially reduced after calibration, implying a strong improvement in reliability indices. CN0 EMOS-Q and CN0 DRN-O forecasts are still slightly underdispersive and biased, showing reliability indices of 0.1718 and 0.1770 (CN0 EMOS-Q) and 0.3357 and 0.3543 (CN0 DRN-Q), whereas the corrected ensemble, as the slightly hump-shaped histograms indicate, is a bit overdispersive (under-confident), especially for short lead times. However, this forecast results in the lowest reliability indices of 0.1380 and 0.0609. Note that the shapes of the histograms of Fig. 8 are completely in line with the characteristics of the corresponding central prediction intervals displayed in Fig. 7. The lower coverage and narrower prediction intervals of the parametric CN0 EMOS and CN0 DRN approaches stem from their underdispersive character, while the mild overdispersion of the corrected ensemble results in wider central prediction intervals and coverage values slightly above the nominal 77.78 % level.

Finally, the MAE of the median of all five investigated predictions (Fig. 9a) and the difference in the MAE of the median from the raw WRF ensemble (Fig. 9b) convey the same message as Fig. 5. Compared to the raw ensemble, post-processing substantially improves the accuracy of the forecast median, and the corrected ensemble results in the lowest MAE values, followed by the DRN and EMOS approaches. Note that the medians of forecasts provided as full predictive distributions (CN0 EMOS and CN0 DRN) differ from the medians of their sample counterparts (CN0 EMOS-Q and CN0 DRN-Q); nevertheless, these differences are tiny as there are hardly any visible dissimilarities in the MAE values of the matching predictions.

The above results indicate that all investigated postprocessing approaches improve the calibration of probabilistic forecasts and the accuracy of point forecasts, and the corrected ensemble exhibits the best overall performance. The CN0 DRN model outperforms its EMOS counterpart in terms of both the mean CRPS and the MAE of the forecast median,

S. Baran et al.: Machine-learning-based probabilistic forecasting of solar irradiance in Chile



Figure 9. (a) MAE of the median of raw and post-processed irradiance forecasts and (b) difference in the MAE of the median forecasts from the raw WRF ensemble as functions of the lead time.

but it is slightly more underdispersive, implying lower coverage and higher reliability indices.

Finally, note that we have found no evidence of a clear dependence of the predictive performance of post-processed forecasts on the location or altitude. Nevertheless, the smallest improvement, for instance, in the mean CRPS compared to the raw ensemble, can be observed at station nos. 17 (El Tololo) and 30 (Liceo Samuel Román Rojas), where the median error of the WRF forecasts is the smallest (see Fig. 2). The former is the only location where the CRPSS of the worst-performing CN0 EMOS appears to be negative.

5 Conclusions

We investigate the skill of probabilistic solar irradiance forecasts for the Atacama and Coquimbo regions in Chile. These areas play a crucial role in photovoltaic power production in the country, and our results can help obtain accurate probabilistic power forecasts. For the sake of this study, the Advanced Research module of the Weather Research and Forecasting (WRF) model was utilized to generate short-term ensemble forecasts of solar irradiance for 30 locations, with forecast horizons ranging up to 48 h and with a temporal resolution of 1 h. The forecasts comprise eight members differing in terms of the planetary boundary layer and radiation parameterization. When verified against station observations, the raw ensemble forecasts exhibit a systematic positive bias for all forecast horizons and are underdispersive; however, the magnitude of the forecast error varies substantially for the different locations.

For post-processing, we consider the parametric EMOS and DRN approaches based on a normal distribution left censored at zero and a distribution-free ensemble correction technique, where a neural network is trained to produce improved ensemble predictions. While all investigated post-processing methods substantially improve the calibration of probabilistic forecasts and reduce the MAE of the forecast median, the latter approach shows the best overall performance. From the competing parametric methods, the machine-learning-based DRN outperforms the corresponding EMOS model, which aligns with the findings of Baran and Baran (2024) and Horat et al. (2025).

As demonstrated by Horat et al. (2025), the calibration of solar irradiance ensemble predictions directly results in enhancements in the calibration of the corresponding probabilistic PV forecasts obtained by applying the PV model chains to each ensemble member separately. They also showed that even better results can be obtained by direct calibration of the PV power ensemble predictions; however, this latter approach requires the availability of the corresponding PV power production data as well, which are rather hard to obtain.

Note that all of the investigated post-processing methods provide a rather general framework and, with slight modifications, adapting them to the specialties of the input data and optimizing the training method, can be applied to any EPS and any geographical region. CN0 EMOS has already proved to be successful in calibrating 40-member ICON-EPS direct and diffuse irradiance forecasts of the DWD for three major cities in Germany (Schulz et al., 2021); 50-member GHI forecasts of the ECMWF IFS for the Jacumba Solar Project in southern California, US (Horat et al., 2025); and 11-member AROME-EPS GHI forecasts for seven solar farms in Hungary (Baran and Baran, 2024), where this post-processing approach is also in its pre-operational testing phase at the HMS. The latter two datasets have also been applied to investigate the skills of various CN0 DRN models, which, as mentioned, consistently outperformed their EMOS counterparts. The only limitation is the availability of training data; one should have enough forecast-observation pairs for reliable model training, especially for the machinelearning-based methods.

Although this first study demonstrates the weaknesses of WRF irradiance ensemble forecasts and the potential in their statistical post-processing, there is still space for further improvements in calibration. On the one hand, one might try to augment the input features of the machine-learning-based forecasts with predictions of related quantities. We have performed tests where WRF ensemble forecasts of temperature and wind speed were also included; however, the gain in forecast skill was minor, if any. On the other hand, one can test advanced neural-network-based distribution-free methods such as the Bernstein quantile network (Bremnes, 2020), where the forecast distribution is modelled with a linear combination of Bernstein polynomials or with the histogram estimation network investigated, for instance, in Scheuerer et al. (2020).

Another interesting direction would be the investigation of various multivariate post-processing methods to produce temporary consistent forecast trajectories. Here, one can think of both the state-of-the-art two-step approaches, such as the ensemble copula coupling (Schefzik et al., 2013) or the Schaake shuffle (Clark et al., 2004), where the CN0 EMOS-Q, the CN0 DRN-Q, or the corrected ensemble can serve as initial independent forecasts, and the recent data-driven conditional generative model of Chen et al. (2024).

Code availability. Python implementations of the neural-network-based post-processing methods and R codes of the CN0 EMOS model are available at https://doi.org/10.5281/zenodo.15612831 (Nagy-Lakatos, 2025).

Data availability. The data investigated in this study are directly available from the authors upon request.

Author contributions. SB: conceptualization, formal analysis, funding acquisition, investigation, methodology, software, validation, visualization, writing (original draft; review and editing). JCM: data curation, methodology, software, visualization, writing (original draft; review and editing). OC: data curation, methodology, software. MD: data curation, methodology, software. MS: data curation, software, writing (review and editing). ON: funding acquisition, writing (review and editing). ML: conceptualization, data curation, investigation, methodology, software, validation, writing (original draft; review and editing).

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors. **Special issue statement.** This article is part of the special issue "Artificial intelligence and machine learning in climate and weather science research". It is not associated with a conference.

Acknowledgements. The authors gratefully acknowledge the support of the S&T cooperation program of the National Research, Development and Innovation Office (grant no. 2021-1.2.4-TÉT-2021-00020). Sándor Baran, Mária Lakatos, and Marianna Szabó were also supported by the National Research, Development and Innovation Office under grant no. K142849. Julio César Marín and Omar Cuevas acknowledge the support of the Center of Atmospheric Studies and Climate Change of the University of Valparaíso, Chile. Last, but not least, the authors thank the two anonymous reviewers, whose constructive comments helped to improve the paper.

Financial support. This research has been supported by the National Research, Development and Innovation Office (grant nos. 2021-1.2.4-TÉT-2021-00020 and K142849).

Review statement. This paper was edited by Lyndsay Shand and reviewed by two anonymous referees.

References

- Bakker, K., Whan, K., Knap, W., and Schmeits, M.: Comparison of statistical post-processing methods for probabilistic NWP forecasts of solar radiation, Sol. Energy, 191, 138–150, https://doi.org/10.1016/j.solener.2019.08.044, 2019.
- Baran, Á. and Baran, S.: A two-step machine learning approach to statistical post-processing of weather forecasts for power generation, Q. J. Roy. Meteorol. Soc., 150, 1029–1047, https://doi.org/10.1002/qj.4635, 2024.
- Baran, S. and Lakatos, M.: Clustering-based spatial interpolation of parametric post-processing models, Weather Forecast., 39, 1591–1604, https://doi.org/10.1175/WAF-D-24-0016.1, 2024.
- Baran, S. and Nemoda, D.: Censored and shifted gamma distribution based EMOS model for probabilistic quantitative precipitation forecasting, Environmetrics, 27, 280–292, https://doi.org/10.1002/env.2391, 2016.
- Baran, S., Baran, Á., Pappenberger, F., and Ben Bouallègue, Z.: Statistical post-processing of heat index ensemble forecasts: is there a royal road? Q. J. Roy. Meteorol. Soc., 146, 3416–3434, https://doi.org/10.1002/qj.3853, 2020.
- Bauer, P., Thorpe, A., and Brunet, G.: The quiet revolution of numerical weather prediction, Nature, 525, 47–55, https://doi.org/10.1038/nature14956, 2015.
- Bi, K., Xie, L., Zhang, H., Chen, X., and Tian, Q.: Accurate medium-range global weather forecasting with 3D neural networks, Nature, 619, 533–538, https://doi.org/10.1038/s41586-023-06185-3, 2023.
- Bremnes, J. B.: Constrained quantile regression splines for ensemble postprocessing, Mon. Weather Rev., 147, 1769–1780, https://doi.org/10.1175/MWR-D-18-0420.1, 2019.

- Bremnes, J. B.: Ensemble postprocessing using quantile function regression based on neural networks and Bernstein polynomials, Mon. Weather Rev., 148, 403–414, https://doi.org/10.1175/MWR-D-19-0227.1, 2020.
- Bretherton, C. S. and Park, S.: A new moist turbulence parameterization in the Community Atmosphere Model, J. Climate, 22, 3422–3448, https://doi.org/10.1175/2008JCLI2556.1, 2009.
- Buizza, R.: Introduction to the special issue on "25 years of ensemble forecasting", Q. J. Roy. Meteorol. Soc., 145, 1–11, https://doi.org/10.1002/qj.3370, 2018a.
- Buizza, R.: Ensemble forecasting and the need for calibration, in: Statistical Postprocessing of Ensemble Forecasts, edited by: Vannitsem, S., Wilks, D. S., and Messner, J. W., Elsevier, Amsterdam, 15–48, https://doi.org/10.1016/B978-0-12-812372-0.00002-9, 2018b.
- Clark, M., Gangopadhyay, S., Hay, L., Rajagopalan, B., and Wilby, R.: The Schaake shuffle: A method for reconstructing space-time variability in forecasted precipitation and temperature fields, J. Hydrometeorol., 5, 243–262, https://doi.org/10.1175/1525-7541(2004)005<0243:TSSAMF>2.0.CO;2, 2004.
- Chen, J., Janke, T., Steinke, F., and Lerch, S.: Generative machine learning methods for multivariate ensemble postprocessing, Ann. Appl. Stat., 18, 159–183, https://doi.org/10.1214/23-AOAS1784, 2024.
- Chou, M.-D. and Suarez, M. J.: A solar radiation parameterization for atmospheric studies, NASA Tech. Memo. 15, NASA/TM-1999-104606, NASA, https://ntrs.nasa.gov/citations/ 19990060930 (last access: 28 February 2025), 1999.
- Chou, M.-D., Suarez, M. J., Liang, X.-Z., and Yan, M. M.-H.: A thermal infrared radiation parameterization for atmospheric studies, NASA Tech. Memo. 19, NASA/TM-2001-104606, NASA, https://ntrs.nasa.gov/citations/20010072848 (last access: 28 February 2025), 2001.
- Delle Monache, L., Hacker, J. P., Zhou, Y., Deng, X., and Stull, R. B.: Probabilistic aspects of meteorological and ozone regional ensemble forecasts, J. Geophys. Res., 111, D24307, https://doi.org/10.1029/2005JD006917, 2006.
- Díaz, M., Nicolis, O., Marín, J. C., and Baran, S.: Statistical post-processing of ensemble forecasts of temperature in Santiago de Chile, Meteorol. Appl., 27, e1818, https://doi.org/10.1002/met.1818, 2020.
- Díaz, M., Nicolis, O., Marín, J. C., and Baran, S.: Postprocessing methods for calibrating the wind speed forecasts in central regions of Chile, Ann. Math. Inform., 53, 93–108, https://doi.org/10.33039/ami.2021.03.012, 2021.
- Dudhia, J.: Numerical study of convection observed Winter Monsoon Experiment during the using mesoscale two-dimensional model, J. а Atmos. 3077-3107, Sci., 46, https://doi.org/10.1175/1520-0469(1989)046<3077:NSOCOD>2.0.CO;2, 1989.
- ECMWF: IFS Documentation CY49R1 Part V: Ensemble Prediction System, ECMWF, Reading, https://doi.org/10.21957/956d60ad81, 2024.
- Friederichs, P. and Hense, A.: Statistical downscaling of extreme precipitation events using censored quantile regression, Mon. Weather Rev., 135, 2365–2378, https://doi.org/10.1175/MWR3403.1, 2007.

- Gneiting, T.: Making and evaluating point forecasts, J. Am. Stat. Assoc., 106, 746–762, https://doi.org/10.1198/jasa.2011.r10138, 2011.
- Gneiting, T. and Raftery, A. E.: Strictly proper scoring rules, prediction and estimation, J. Am. Stat. Assoc., 102, 359–378, https://doi.org/10.1198/016214506000001437, 2007.
- Gneiting, T., Raftery, A. E., Westveld, A. H., and Goldman, T.: Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation, Mon. Weather Rev., 133, 1098–1118, https://doi.org/10.1175/MWR2904.1, 2005.
- Goodfellow, I., Bengio, Y., and Courville, A.: Deep Learning, MIT Press, Cambridge, ISBN 978-0262035613, 2016.
- Gu, Y., Liou, K. N., Ou, S. C., and Fovell, R.: Cirrus cloud simulations using WRF with improved radiation parameterization and increased vertical resolution, J. Geophys. Res., 116, D06119, https://doi.org/10.1029/2010JD014574, 2011.
- Hamill, T. M. and Scheuerer, M.: Probabilistic precipitation forecast postprocessing using quantile mapping and rank-weighted best-member dressing, Mon. Weather Rev., 146, 4079–4098, https://doi.org/10.1175/MWR-D-18-0147.1, 2018.
- Hong, S.-Y., Noh, Y., and Dudhia, J.: A new vertical diffusion package with an explicit treatment of entrainment processes, Mon. Weather Rev., 134, 2318–2341, https://doi.org/10.1175/MWR3199.1, 2006.
- Horat, N. Klerings, S., and Lerch, S.: Improving model chain approaches for probabilistic solar energy forecasting through postprocessing and machine learning, Adv. Atmos. Sci., 42, 297–312, https://doi.org/10.1007/s00376-024-4219-2, 2025.
- Hu, Y., Schmeits, M. J., van Andel, J. S., Verkade, J. S., Xu, M., Solomatine, D. P., and Liang, Z.: A stratified sampling approach for improved sampling from a calibrated ensemble forecast distribution, J. Hydrometeorol., 17, 2405–2417, https://doi.org/10.1175/JHM-D-15-0205.1, 2016.
- Iacono, M. J., Delamere J. S., Mlawer, E. J., Shephard, M. W., Clough, S. A., and Collins, W. D.: Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models, J. Geophys. Res., 113, D13103, https://doi.org/10.1029/2008JD009944, 2005.
- IRENA: Renewable capacity statistics 2024, International Renewable Energy Agency, Abu Dhabi, ISBN 978-92-9260-587-2, 2024.
- Janjić, Z. I.: The Step-Mountain Eta Coordinate Model: further developments of the convection, viscous sublayer, and turbulence closure schemes, Mon. Weather Rev., 122, 927–945, https://doi.org/10.1175/1520-0493(1994)122<0927:TSMECM>2.0.CO;2, 1994.
- Jávorné Radnóczi, K., Várkonyi, A., and Szépszó, G.: On the way towards the AROME nowcasting system in Hungary, ALADIN-HIRLAM Newsletter, 65–69, https://www.umr-cnrm.fr/aladin/ IMG/pdf/nl14.pdf (last access: 28 February 2025), 2020.
- Jobst, D., Möller, A., and Groß, J.: D-vine-copula-based postprocessing of wind speed ensemble forecasts, Q. J. Roy. Meteorol. Soc., 149, 2575–2597, https://doi.org/10.1002/qj.4521, 2023.
- Jordan, A., Krüger, F., and Lerch, S.: Evaluating probabilistic forecasts with scoringRules, J. Stat. Softw., 90, 1–37, https://doi.org/10.18637/jss.v090.i12, 2019.
- Kain, J. S.: The Kain-Fritsch convective parameterization: an update, J. Appl. Meteo-

rol., 43, 170–181, https://doi.org/10.1175/1520-0450(2004)043<0170:TKCPAU>2.0.CO;2, 2004.

- Krüger, F., Lerch, S., Thorarinsdottir, T. L., and Gneiting, T.: Predictive inference based on Markov chain Monte Carlo output, Int. Stat. Rev., 89, 274–301, https://doi.org/10.1111/insr.12405, 2021.
- Lang, M. N., Lerch, S., Mayr, G. J., Simon, T., Stauffer, R., and Zeileis, A.: Remember the past: a comparison of time-adaptive training schemes for non-homogeneous regression, Nonlin. Processes Geophys., 27, 23–34, https://doi.org/10.5194/npg-27-23-2020, 2020.
- Lang, S., Alexe, M., Chantry, M., Dramsch, J., Pinault, F., Raoult, B., Ben Bouallègue, Z., Clare, M., Lessig, C., Magnusson, L., and Prieto Nemesio, A.: AIFS: a new ECMWF forecasting system, ECMWF Newsletter, 4–5, https://www.ecmwf.int/ en/newsletter/178 (last access: 28 February 2025), 2024.
- La Salle, J. L. G., Badosa, J., David, M., Pinson, P., and Lauret, P.: Added-value of ensemble prediction system on the quality of solar irradiance probabilistic forecasts, Renew. Energy, 162, 1321–1339, https://doi.org/10.1016/j.renene.2020.07.042, 2020.
- Lerch, S. and Baran, S.: Similarity-based semi-local estimation of EMOS models, J. Roy. Stat. Soc. Ser. C, 66, 29–51, https://doi.org/10.1111/rssc.12153, 2017.
- Li, W., Pan, B., Xia, J., and Duan, Q.: Convolutional neural network-based statistical post-processing of ensemble precipitation forecasts, J. Hydrol., 605, 127301, https://doi.org/10.1016/j.jhydrol.2021.127301, 2022.
- Mayer, M. J. and Yang, D.: Probabilistic photovoltaic power forecasting using a calibrated ensemble of model chains, Renew. Sustain. Energ. Rev., 168, 112821, https://doi.org/10.1016/j.rser.2022.112821, 2022.
- Molina, A., Falvey, M., and Rondanelli, R.: A solar radiation database for Chile, Sci. Rep., 7, 14823, https://doi.org/10.1038/s41598-017-13761-x, 2017.
- Nagy-Lakatos, M.: marialakatos/pp_radiation_forecasts: Initial release (v1.0.0), Zenodo [code], https://doi.org/10.5281/zenodo.15612831, 2025.
- Niu, G.-Y., Yang, Z.-L., Mitchell, K. E., Chen, F., Ek, M. B., Barlage, M., Kumar, A., Manning, K., Niyogi, D., Rosero, E., Tewari, M., and Xia, Y.: The community Noah land surface model with multiparameterization options (Noah-MP): 1. Model description and evaluation with local-scale measurements, J. Geophys. Res., 116, D12109, https://doi.org/10.1029/2010JD015139, 2011.
- Pleim, J. E.: A combined local and nonlocal closure model for the atmospheric boundary layer. Part I: Model description and testing, J. Appl. Meteorol. Clim., 46, 1383–1395, https://doi.org/10.1175/JAM2539.1, 2007.
- Politis, D. N. and Romano, J. P.: The stationary bootstrap, J. Am. Stat. Assoc., 89, 1303–1313, https://doi.org/10.2307/2290993, 1994.
- Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T. R., El-Kadi, A., Masters, D., Ewalds, T., Stott, J., Mohamed, S., Battaglia, P., Lam, R., and Willson, M.: Probabilistic weather forecasting with machine learning, Nature, 637, 84–90, https://doi.org/10.1038/s41586-024-08252-9, 2025.
- Rasp, S. and Lerch, S.: Neural networks for postprocessing ensemble weather forecasts, Mon. Weather Rev., 146, 3885–3900, https://doi.org/10.1175/MWR-D-18-0187.1, 2018.

- Rondanelli, R., Molina, A., and Falvey, M.: The Atacama surface solar maximum, B. Am. Meteorol. Soc., 96, 405–418, https://doi.org/10.1175/BAMS-D-13-00175.1, 2015.
- Schefzik, R., Thorarinsdottir, T. L., and Gneiting, T.: Uncertainty quantification in complex simulation models using ensemble copula coupling, Stat. Sci., 28, 616–640, https://doi.org/10.1214/13-STS443, 2013.
- Scheuerer, M.: Probabilistic quantitative precipitation forecasting using ensemble model output statistics, Q. J. Roy. Meteorol. Soc., 140, 1086–1096, https://doi.org/10.1002/qj.2183, 2014.
- Scheuerer, M., Switanek, M. B., Worsnop, R. P., and Hamill, T. M.: Using artificial neural networks for generating probabilistic subseasonal precipitation forecasts over California, Mon. Weather Rev., 148, 3489–3506, https://doi.org/10.1175/MWR-D-20-0096.s1, 2024.
- Schulz, B. and Lerch, S.: Machine learning methods for postprocessing ensemble forecasts of wind gusts: a systematic comparison, Mon. Weather Rev., 150, 235–257, https://doi.org/10.1175/MWR-D-21-0150.1, 2022.
- Schulz, B., El Ayari, M., Lerch, S., and Baran, S.: Postprocessing numerical weather prediction ensembles for probabilistic solar irradiance forecasting, Sol. Energy, 220, 1016– 1031, https://doi.org/10.1016/j.solener.2021.03.023, 2021.
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Liu, Z., Berner, J., Wang, W., Powers, J. G., Duda, M. G., Barker, D. M., and Huang, X. Y.: A description of the advanced research WRF version 4, NCAR Tech. Note NCAR/TN-556+STR, NCAR, https://doi.org/10.5065/1DFH-6P97, 2019.
- Song, M., Yang, D., Lerch, S., Xia, X., Yagli, G. M., Bright, J. M., Shen, Y., Liu, B., Liu, X., and Mayer, M. J.: Non-crossing quantile regression neural network as a calibration tool for ensemble weather forecasts, Adv. Atmos. Sci., 41, 1417–1437, https://doi.org/10.1007/s00376-023-3184-5, 2024.
- Sukoriansky, S., Galperin, B., and Perov, V.: Application of a new spectral theory of stably stratified turbulence to the atmospheric boundary layer over sea ice, Bound.-Lay. Meteorol., 117, 231– 257, https://doi.org/10.1007/s10546-004-6848-4, 2005.
- Szabó, M., Gascón, E., and Baran, S.: Parametric post-processing of dual-resolution precipitation forecasts, Weather Forecast., 38, 1313–1322, https://doi.org/10.1175/WAF-D-23-0003.1, 2023.
- Taillardat, M., Mestre, O., Zamo, M., and Naveau, P.: Calibrated ensemble forecasts using quantile regression forests and ensemble model output statistics, Mon. Weather Rev., 144, 2375–2393, https://doi.org/10.1175/MWR-D-15-0260.1, 2016.
- Thompson, G., Field, P. R., Rasmussen, R. M., and Hall, W. D.: Explicit forecasts of winter precipitation using an improved bulk microphysics scheme. Part II: Implementation of a new snow parameterization, Mon. Weather Rev., 136, 5095–5115, https://doi.org/10.1175/2008MWR2387.1, 2008.
- Thorarinsdottir, T. L. and Gneiting, T.: Probabilistic forecasts of wind speed: ensemble model output statistics by using heteroscedastic censored regression, J. Roy. Stat. Soc. Ser. A, 173, 371–388, https://doi.org/10.1111/j.1467-985X.2009.00616.x, 2010.
- Vannitsem, S., Bremnes, J. B., Demaeyer, J., Evans, G. R., Flowerdew, J., Hemri, S., Lerch, S., Roberts, N., Theis, S., Atencia, A., Ben Boualègue, Z., Bhend, J., Dabernig, M., De Cruz, L., Hieta, L., Mestre, O., Moret, L., Odak Plenkovič, I., Schmeits, M., Taillardat, M., Van den Bergh, J., Van Schaey-

broeck, B., Whan, K., and Ylhaisi, J.: Statistical postprocessing for weather forecasts – review, challenges and avenues in a big data world, B. Am. Meteorol. Soc., 102, E681–E699, https://doi.org/10.1175/BAMS-D-19-0308.1, 2021.

- Van Schaeybroeck, B. and Vannitsem, S.: Ensemble postprocessing using member-by-member approaches: Theoretical aspects, Q. J. Roy. Meteorol. Soc., 141, 807–818, https://doi.org/10.1002/qj.2397, 2015.
- Veldkamp, S., Whan, K., Dirksen, S., and Schmeits, M.: Statistical postprocessing of wind speed forecasts using convolutional neural networks, Mon. Weather Rev., 149, 1141–1152, https://doi.org/10.1175/MWR-D-20-0219.1, 2021.
- Wilks, D. S.: Statistical Methods in the Atmospheric Sciences, in: 4th Edn., Elsevier, Amsterdam, ISBN 978-0-12-815823-4, https://doi.org/10.1016/C2017-0-03921-6, 2019.
- Yang, D.: Ensemble model output statistics as a probabilistic siteadaptation tool for solar irradiance: A revisit, J. Renew. Sustain. Energ., 12, 036101, https://doi.org/10.1063/5.0010003, 2020.

- Yang, D. and Kleissl, J.: Solar Irradiance and Photovoltaic Power Forecasting, CRC Press, Boca Raton, https://doi.org/10.1201/9781003203971, 2024.
- Yang, D. and van der Meer, D.: Post-processing in solar forecasting: Ten overarching thinking tools, Renew. Sustain. Energ. Rev., 140, 110735, https://doi.org/10.1016/j.rser.2021.110735, 2021.
- Yang, Z.-L., Niu, G.-Y., Mitchell, K. E., Chen, F., Ek, M. B., Barlage, M., Longuevergne, L., Manning, K., Niyogi, D., Tewari, M., and Xia, Y.: The community Noah land surface model with multiparameterization options (Noah-MP): 2. Evaluation over global river basins, J. Geophys. Res., 116, D12110, https://doi.org/10.1029/2010JD015140, 2011.
- Zängl, G., Reinert, D., Rípodas, P., and Baldauf, M.: The ICON (ICOsahedral Non-hydrostatic) modelling framework of DWD and MPI-M: Description of the non-hydrostatic dynamical core, Q. J. Roy. Meteorol. Soc., 141, 563–579, https://doi.org/10.1002/qj.2378, 2015.