

Post-processing numerical weather prediction ensembles for probabilistic solar irradiance forecasting

Benedikt Schulz^a, Mehrez El Ayari^{b,c}, Sebastian Lerch^{a,d}, Sándor Baran^{b,*}

^a Institute for Stochastics, Karlsruhe Institute of Technology, Englerstraße 2, D-76128 Karlsruhe, Germany

^b Faculty of Informatics, University of Debrecen, Kassai út 26, H-4028 Debrecen, Hungary

^c Doctoral School of Informatics, University of Debrecen, Kassai út 26, H-4028 Debrecen, Hungary

^d Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

ARTICLE INFO

Keywords:

Energy forecasting
Ensemble model output statistics
Ensemble post-processing
Probabilistic forecasting
Solar energy
Solar irradiance

ABSTRACT

In order to enable the transition towards renewable energy sources, probabilistic energy forecasting is of critical importance for incorporating volatile power sources such as solar energy into the electrical grid. Solar energy forecasting methods often aim to provide probabilistic predictions of solar irradiance. In particular, many hybrid approaches combine physical information from numerical weather prediction models with statistical methods. Even though the physical models can provide useful information at intra-day and day-ahead forecast horizons, ensemble weather forecasts from multiple model runs are often not calibrated and show systematic biases. We propose a post-processing model for ensemble weather predictions of solar irradiance at temporal resolutions between 30 min and 6 h. The proposed models provide probabilistic forecasts in the form of a censored logistic probability distribution for lead times up to 5 days and are evaluated in two case studies covering distinct physical models, geographical regions, temporal resolutions, and types of solar irradiance. We find that post-processing consistently and significantly improves the forecast performance of the ensemble predictions for lead times up to at least 48 h and is well able to correct the systematic lack of calibration.

1. Introduction

To reduce emissions of greenhouse gases a transition towards renewable energy sources such as wind and solar power is imperative (Van der Meer et al., 2018). Accurate and reliable forecasts of power generation from those sources are thus becoming increasingly important for integrating volatile power systems into the electrical grid in order to balance demand and supply (Gottwalt et al., 2016; Ordiano et al., 2020). The literature on energy forecasting has primarily focused on deterministic prediction for the past decades. However, it has now been widely argued that probabilistic forecasting is essential for optimal decision making in planning and operation (Hong and Fan, 2016; Van der Meer et al., 2018; Haupt et al., 2019; Hong et al., 2020). For example, Hong et al. (2020) identify probabilistic forecasting with the aim of providing a predictive probability distribution for a future quantity or event in order to quantify forecast uncertainty as one of the most important emerging research topics in their recent review on energy forecasting.

We here focus on solar energy which is one of the most important

sources of renewable energy. For example, photovoltaic (PV) power contributes significantly to the power supply in Germany and generated 8.2% of the gross electricity consumption in 2019, and temporarily up to 50% of the current electricity consumption on sunny days (Fraunhofer Institute for Solar Energy Systems, 2020).

Solar energy forecasting approaches can be distinguished into those that aim to predict solar irradiance, and those that aim to predict PV power. Naturally, solar irradiance and PV system output are strongly correlated, and the employed statistical methods are similar (Van der Meer et al., 2018). We will focus on probabilistic solar irradiance forecasting in the following.

For recent comprehensive overviews and reviews of existing approaches, see Van der Meer et al. (2018) and Yang (2019). Except for short-term prediction (e.g., Zelikman et al., 2020), most methods for probabilistic solar irradiance forecasting combine physical information from numerical weather prediction (NWP) models with statistical methods. NWP models describe atmospheric processes via systems of partial differential equations and are often run several times with varying initial conditions and model physics, resulting in an ensemble of

* Corresponding author.

E-mail address: baran.sandor@inf.unideb.hu (S. Baran).

<https://doi.org/10.1016/j.solener.2021.03.023>

Received 16 January 2021; Received in revised form 2 March 2021; Accepted 8 March 2021

0038-092X/© 2021 The Author(s). Published by Elsevier Ltd on behalf of International Solar Energy Society. This is an open access article under the CC

BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

predictions that provide a probabilistic forecast (Bauer et al., 2015). Despite significant advances over the past decades ensemble forecasts continue to exhibit systematic errors that require correction. To that end, a rapidly growing collection of methods for statistical post-processing has been developed in the meteorological and statistical literature. Post-processing methods can be viewed as distributional regression models that produce probabilistic forecasts of a target variable (such as solar irradiance), using ensemble predictions of (potentially many) meteorological variables as input. Prominent methods for post-processing include the ensemble model output statistics (EMOS) approach proposed in Gneiting et al. (2005), where the forecast distribution is given by a parametric probability distribution with parameters connected to the ensemble predictions of the variable of interest via link functions. Other widely used post-processing techniques build on quantile regression methods which provide nonparametric probabilistic forecasts in the form of quantiles of the forecast distribution. We refer to Wilks (2018) for a general introduction and to Vannitsem et al. (2021) for a comprehensive overview of current research developments and operational implementations.

Post-processing ensemble predictions of solar irradiance has recently received a growing interest in the literature on probabilistic solar energy forecasting.¹ Bakker et al. (2019) compare post-processing methods for the clear-sky index that take deterministic NWP predictions of several variables as input, and employ various machine learning approaches for quantile regression. La Salle et al. (2020) propose an EMOS model for global horizontal irradiance where truncated normal and generalized extreme value distributions are used as forecast distributions, and compare to quantile regression and analog ensemble methods. Yaghi et al. (2020) compare several parametric and nonparametric post-processing methods for hourly clear-sky index forecasting, including EMOS models based on Gaussian, truncated logistic and skewed Student's *t* distributions as well as quantile regression based on random forests, and generalized additive models for location, scale and shape. In closely related work, Yang (2020c) proposes the use of EMOS models for probabilistic site adaptation of gridded solar irradiance products, and Yang (2020d) compares building models for irradiance and for the clear-sky index and investigates the choice of parametric distributions. A comprehensive overview of post-processing methods in solar forecasting is provided in the recent review paper of Yang and van der Meer (2021).

Here, we build on the EMOS framework and propose post-processing models for different measured types of solar irradiance, using corresponding NWP ensemble predictions as input. Within the classification introduced in Yang and van der Meer (2021), this corresponds to a *probabilistic-to-probabilistic* post-processing method. To account for the specific discrete-continuous nature of solar irradiance due to the positive probability of observing zero irradiance during night-time, we use a censored logistic forecast distribution motivated by similar models for post-processing ensemble forecasts of precipitation accumulation (Scheuerer, 2014; Baran and Nemoda, 2016). The post-processing models are applied in two case studies that focus on distinct solar irradiance variables, NWP models, temporal resolutions and geographic regions (Hungary and Germany), for lead times of up to 48 and 120 hours, respectively. We utilize periodic models to better capture seasonal variation in solar irradiance, and investigate different temporal compositions of training datasets for model estimation. Further, we compare the effects using a clear-sky index as target variable to stationarize the time series of irradiances – the standard approach in solar forecasting – and post-processing the irradiance forecasts directly.

The remainder of the paper is organized as follows. Section 2 introduces the datasets used in the case studies. The proposed post-

processing approach, training procedures and forecast evaluation methods are described in Section 3, with results for the two case studies presented in Section 4. The paper concludes with a discussion in Section 5, additional results are deferred to the Appendix.

2. Data

In the case studies of Section 4, we evaluate the EMOS models proposed in Section 3.2 using forecasts of different types of solar irradiance produced by two different ensemble prediction systems (EPSs), which cover distinct forecast domains.

2.1. AROME-EPS

The 11-member Applications of Research to Operations at Mesoscale EPS (AROME-EPS) of the Hungarian Meteorological Service (HMS) covers the Transcarpathian Basin with a horizontal resolution of 2.5 km (Jávoriné Radnóczy et al., 2020). It consists of 10 ensemble members obtained from perturbed initial conditions and a control member from an unperturbed analysis. The dataset at hand contains ensemble forecasts of instantaneous values of global horizontal irradiance (GHI) (W/m^2) together with the corresponding validation observations of the HMS for seven representative locations in Hungary (Aszód, Budapest, Debrecen, Kecskemét, Pécs, Szeged, Tápíószele) for the period between 7 May 2020 and 14 October 2020. Forecasts are initialized at 00 UTC with a forecast horizon of 48h and a temporal resolution of 30 minutes resulting in a total of 96 forecasts per submission. Rephrasing this in the terminology of Yang (2020b), the temporal information can be written as

$$\{S^{48h}, \mathcal{R}_f^{30min}, \mathcal{L}^{0h}, \mathcal{U}^{24h}\} \quad \text{and} \quad \mathcal{R}_f = \mathcal{R}_x,$$

where S refers to the forecast span covered by all forecasts, \mathcal{R}_f to the temporal resolution of the forecasts, \mathcal{R}_x to that of the observations, \mathcal{L} to the time between submission and first time stamp, and \mathcal{U} to the time interval between consecutive submissions.

For the AROME-EPS, we will refer to the term lead time as the time between the initialization and the time stamp of the corresponding instantaneous forecast.

2.2. ICON-EPS

The 40-member global ICOSahedralNonhydrostatic EPS (ICON-EPS; Zängl et al., 2015) of the German Meteorological Service (DWD; Deutsche Wetterdienst) was launched in 2018 and has a horizontal resolution of 20 km over Europe (ICON-EU EPS). The ensemble members are generated with the help of random perturbations and the forecasts are initialized four times a day at 00/06/12/18 UTC each with a forecast horizon of 120h (Reinert et al., 2020). Within the first 48h the temporal resolution is 1h, 3-hourly up to 72h, while between 72 and 120h the resolution is 6h (a total of 64 forecasts). Adapting the terminology of Yang (2020b) accordingly, this corresponds to

$$\{S^{(48h,24h,48h)}, \mathcal{R}_f^{(1h,3h,6h)}, \mathcal{L}^{(0h,48h,72h)}, \mathcal{U}^{6h}\}.$$

The ICON ensemble predictions are given as averages between two time stamps, e.g. the 12-step ahead forecast is the average predicted irradiance between 11 to 12h after initialization time and the 59-step ahead forecast is the average from 84 to 90h. For simplicity, we will refer to a individual forecast by the lead time and not the step ahead, where lead time refers to the time between submission and the final time stamp, i.e., the former forecast has a lead time of 12h, the latter a lead time of 90h.

Our dataset contains ensemble forecasts of the two components of GHI: beam normal irradiance (BNI) adjusted for the solar zenith angle θ (i.e., $\text{BNI} \cdot \cos(\theta)$), and diffuse horizontal radiation (DHI) (W/m^2). To

¹ Note that related post-processing methods have also been applied for direct forecasting of PV power output, for example in Sperati et al. (2016). We limit our literature review to methods that aim to predict solar irradiance to retain focus.




Fig. 1. Illustration of the (generalized) PDF and CDF of a logistic distribution with parameters $\mu = -0.5$ and $\sigma = 0.5$, and the effects of truncating and left-censoring at 0.

simplify the distinction between the different types of irradiance and improve the readability of the article, we will refer to $\text{BNI}\cdot\cos(\theta)$ as beam horizontal irradiance (BHI) or direct irradiance, to DHI as diffuse irradiance, and to $\text{GHI} = \text{BHI} + \text{DHI}$ as global irradiance, in the following.

We further obtained corresponding observational data for weather stations located near the major cities of Berlin, Hamburg and Karlsruhe from Open Data Server of DWD (DWD Climate Data Center, 2020). The observations are computed based on 10-minute sums of the corresponding variables, i.e. the temporal resolution of the observation is $\mathcal{R}_x^{10\text{min}}$. For detailed descriptions, we refer to Reinert et al. (2020) for the ensemble predictions and Becker and Behrens (2012) for the observations. The entire dataset used here covers the period 27 December 2018 – 31 December 2020.

3. Post-processing methods and forecast evaluation

Broadly, two kinds of statistical post-processing methods can be distinguished. Parametric (or distribution-based) methods specify the distribution of the variable of interest, whereas nonparametric (or distribution-free) methods do not assume a certain forecast distribution and typically forecast a set of quantiles or random draws that represent the distribution. While distribution-free methods do not require the choice of a parametric family of probability distributions for the forecast model, they usually do not offer a complete representation of the forecast distribution. For example, quantile regression approaches only provide predictions for a set of quantiles, and it might not be clear how to extrapolate beyond the range of the considered quantile levels. By contrast, parametric methods provide full predictive distributions from which arbitrary quantities, e.g. quantiles, can be computed. In addition, parametric methods are generally better suited for situations in which only limited amounts of training data are available. In view of the sparsity of data, we here choose the distribution-based EMOS approach. In the following, we propose a novel EMOS model based on a censored logistic distribution.

3.1. Choice of forecast distribution

The discrete-continuous nature of solar irradiance calls for non-negative predictive distributions assigning positive mass to the event of zero irradiance. Similar to parametric approaches to post-processing ensemble forecasts of precipitation accumulation, one can either left-

truncate an appropriate continuous distribution at zero (see e.g., Scheuerer, 2014; Baran and Nemoda, 2016), or choose the more complex method of mixing a point mass at zero and a suitable continuous distribution with non-negative support (Sloughter et al., 2007; Bentzen and Friederichs, 2012). Here, we focus on the former and introduce the censored logistic distribution on which our approach is built.

Consider a logistic distribution $\mathcal{L}(\mu, \sigma)$ with location μ and scale $\sigma > 0$ specified by the probability density function (PDF)

$$g(x; \mu, \sigma) := \frac{e^{-(x-\mu)/\sigma}}{\sigma(1 + e^{-(x-\mu)/\sigma})^2}, \quad x \in \mathbb{R},$$

and the cumulative distribution function (CDF) $G(x; \mu, \sigma) := (1 + e^{-(x-\mu)/\sigma})^{-1}$.

The logistic distribution left-censored at zero (CLO) assigns point mass $G(0; \mu, \sigma) = (1 + e^{\mu/\sigma})^{-1}$ to the origin, i.e. the probability of observing a negative value (before censoring) is the probability of observing zero afterwards. The CLO-distribution can be defined by the CDF

$$G_0^c(x; \mu, \sigma) := \begin{cases} G(x; \mu, \sigma), & x \geq 0, \\ 0, & x < 0. \end{cases} \quad (1)$$

or the generalized PDF

$$g_0^c(x; \mu, \sigma) = 1_{\{x=0\}} G(0; \mu, \sigma) + 1_{\{x>0\}} g(x; \mu, \sigma), \quad (2)$$

where 1_A denotes the indicator function of a set A . The p -quantile q_p ($0 < p < 1$) of (1) equals 0 if $p \leq G(0; \mu, \sigma)$ and $q_p = \mu - \sigma(\log(1-p) - \log p)$, otherwise. With the help of (2) it is straightforward to show that the corresponding mean equals

$$\mu_0^c := \mu + \sigma \log(1 + e^{-\mu/\sigma}).$$

Note that initial tests with a censored normal predictive distribution were performed for the ICON-EPS dataset; however, the results suggested that the proposed CLO-EMOS approach results in slightly improved predictive performance. The choice of parametric families for the forecast distribution has been an important aspect in post-processing research. For considerations in the context of solar irradiance forecasting, see e.g. Yang (2020d), Yagli et al. (2020), and La Salle et al. (2020). In those previous works, forecast distributions are truncated at zero, i.e., probability mass belonging to negative values before

truncation is re-assigned to positive values by restricting the support to the positive half-axis and multiplying the PDF with a constant factor. The key difference of left-censoring at zero is given by the positive probability assigned to the event of observing exactly zero irradiance. Fig. 1 illustrates this difference for a truncated and left-censored logistic distribution.

Our choice of a left-censored distribution was motivated by the aim to obtain a single distributional model that can be applied to all times of day and is able to account for cases where the observation and all or most of the ensemble member predictions are zero, which makes it unnecessary to remove night-time irradiance data during training and inference. Therefore, there is no need to select location- and season-specific times of day that define periods of time where the post-processing model can be applied, which is the case for models based on truncated distributions. In addition, we have observed occasional cases where zero irradiance is observed, but some of the ensemble members predict non-zero values. A model based on a censored distribution is able to correct those deficiencies and might have advantages for applications in automated procedures where post-processing constitutes one of the components and post-processed forecasts serve as inputs for additional modeling steps.

Generally, it has been noted that a single parametric distribution may not result in a perfect fit, and more involved approaches based on mixtures or combinations of several forecast distributions that have been proposed in the meteorological literature might be able to improve performance, but increase model complexity and computational costs (Baran and Lerch, 2016; Baran and Lerch, 2018).

3.2. Ensemble model output statistics models for solar irradiance forecasting

Now, denote by f_1, f_2, \dots, f_K the ensemble member forecasts of solar irradiance for a given location, time point and forecast horizon. In the simplest proposed EMOS model, the location parameter μ and the scale parameter σ of the CLO-distribution are connected to the ensemble members via link functions

$$\mu = \alpha_0 + \alpha_1 f_1 + \dots + \alpha_K f_K + \nu p_0 \quad \text{and} \quad (3)$$

$$\sigma = \exp(\beta_0 + \beta_1 \log S^2),$$

where p_0 and S^2 are the proportion of zero observations and the ensemble variance, respectively, that is

$$p_0 := \frac{1}{K} \sum_{k=1}^K 1_{\{f_k=0\}} \quad \text{and} \quad S^2 := \frac{1}{K-1} \sum_{k=1}^K (f_k - \bar{f})^2,$$

with \bar{f} denoting the ensemble mean. The EMOS coefficients $\alpha_0, \alpha_1, \dots, \alpha_K, \nu$ and β_0, β_1 are estimated according to the optimum score principle of Gneiting and Raftery (2007), i.e., by optimizing the mean value of an appropriate verification metric over the training data consisting of past pairs of forecasts and observations for a given time period.

In order to capture the seasonal variation in solar irradiance, following the ideas of Hemri et al. (2014), we further fit separate periodic models to both observations and ensemble forecasts of the training data. Two regression models dealing with oscillations of a single and two different frequencies are investigated, namely

$$y_t = c_0 + c_1 \sin\left(\frac{2\pi t}{365}\right) + c_2 \cos\left(\frac{2\pi t}{365}\right) + \varepsilon_t \quad \text{and} \quad (4)$$

$$y_t = d_0 + d_1 \sin\left(\frac{2\pi t}{365}\right) + d_2 \cos\left(\frac{2\pi t}{365}\right) + d_3 \sin\left(\frac{4\pi t}{365}\right) + d_4 \cos\left(\frac{4\pi t}{365}\right) + \varepsilon_t, \quad (5)$$

where the dependent variables y_t , $t = 1, 2, \dots, n$, are either irradiance observations for a given location or members of the corresponding ensemble forecast with a given lead time h from a training period of length n . With the help of either (4) or (5) one can calculate the h ahead predictions \hat{y} and \hat{f}_k of the observation and ensemble members, respectively, and consider the following modified link function for the location:

$$\mu = \hat{y} + \alpha_0 + \alpha_1 (f_1 - \hat{f}_1) + \dots + \alpha_K (f_K - \hat{f}_K) + \nu p_0. \quad (6)$$

Model formulations (3) and (6) are valid under the assumption that each ensemble member can be identified and tracked. However, most operationally used EPSs today generate ensemble forecasts that lack individually distinguishable physical features such as distinct variations in the model physics, for example by generating ensemble member based on random perturbations of initial conditions. Those statistically indistinguishable members (or groups of members) generated in this way are usually referred to as exchangeable (Fraley et al., 2010) in reference to the concept of exchangeable random variables in statistics. This is also the case for the ICON-EPS and the AROME-EPS described in Section 2. The existence of groups of exchangeable ensemble members should be taken into account during model formulation. This is usually achieved by requiring that ensemble members within a given group share the same coefficients (see e.g. Wilks, 2018). If there exist M ensemble members divided into K exchangeable groups and \bar{f}_k denotes the mean of the k th group containing M_k ensemble members ($\sum_{k=1}^K M_k = M$), the exchangeable versions of link functions (3) and (6) are

$$\mu = \alpha_0 + \alpha_1 \bar{f}_1 + \dots + \alpha_K \bar{f}_K + \nu p_0 \quad (7)$$

and

$$\mu = \hat{y} + \alpha_0 + \alpha_1 (\bar{f}_1 - \hat{f}_1) + \dots + \alpha_K (\bar{f}_K - \hat{f}_K) + \nu p_0, \quad (8)$$

respectively, where \hat{f}_k is the prediction of \bar{f}_k for lead time h based either on (4) or (5).

3.3. Training data selection

The parameters of the CLO-EMOS model are estimated with the help of ensemble forecasts and corresponding observations from a training data set, where several options in terms of both spatial and temporal composition can be considered. From the spatial point of view, there are two traditional approaches: local and regional (sometimes also called global) selection (Thorarindottir and Gneiting, 2010). In the local approach, the parameters of the predictive distribution for a given location are estimated using only data from that particular location, resulting in different parameter estimates for the different locations. In order to ensure numerical stability of the estimation process, local modeling requires long time periods for training, which is the major disadvantage of this approach. As it addresses the location-specific forecast error characteristics, it often results in better forecast skill than regional estimation, where training data of the whole ensemble domain are used and all locations share the same set of parameters.

Regarding the temporal composition, the standard approach in EMOS modeling is the use of rolling training periods, where training data consists of forecasts and observations for the n calendar days preceding the target date of interest. Rolling training periods can be flexibly applied to smaller datasets and enable models to adapt to changes in meteorological conditions or the underlying NWP system. An alternative approach is to utilize all available data by considering expanding training periods, motivated by studies suggesting that using long archives of training data irrespective of potential NWP model changes during that period often show superior performance (Lang et al., 2020). Regularly extending training sets may for example be relevant in operational implementations where data archives might be built up and expanded over time. In the case studies of Section 4, examples of all listed training data selection methods are shown: regional estimation with a rolling training period in Section 4.1 and local estimation with rolling and extending training periods in Section 4.2.

For both datasets, ensemble predictions of multiple lead times are available, and in case of the ICON-EPS dataset the forecasts are initialized by the NWP model at four different times of the day. These are treated separately when estimating model parameters, i.e., a separate post-processing model is estimated for each lead time and each initialization hour, based on training datasets comprised of data from those lead times and initialization hours only. Thereby, we aim to account for changes in the forecast error characteristics of the raw ensemble predictions over multiple lead times, and for potential diurnal effects by ensuring that the training data covers the same time of day of the observation. Since we do not remove night-time data during training and inference, this further helps to account for positive probabilities of observing zero irradiance as point masses in the forecast distributions. Note that seasonal variations for a given time of day, for example effects of differing solar zenith angles, are implicitly modeled when using (4) or (5).

3.4. Forecast evaluation

As mentioned in Section 3.2, the estimates of the unknown parameters of the CL0-EMOS model minimize the mean of a proper verification score over the training data. These so-called proper scoring rules (Gneiting and Raftery, 2007) have been widely used in the meteorological and economic literature, and have recently become popular tools for evaluating probabilistic solar energy forecasts. See, for example, the recommendations in Lauret et al. (2019) which we will follow throughout.

The most popular proper scoring rule in the atmospheric sciences is the continuous ranked probability score (CRPS, Section 9.5.1 Wilks, 2019), which simultaneously assesses both calibration and sharpness of the probabilistic forecasts. The former refers to a statistical consistency between forecasts and observations, whereas the latter refers to the concentration of the predictive distribution, see Lauret et al. (2019) for details. For a predictive CDF F and real-valued observation x , the CRPS is defined as

$$\begin{aligned} \text{CRPS}(F, x) &:= \int_{-\infty}^{\infty} (F(y) - 1_{\{y \geq x\}})^2 dy \\ &= E|X - x| - \frac{1}{2}E|X - X'|, \end{aligned} \quad (9)$$

where X and X' are independent random variables with CDF F and finite first moment. The CRPS is a negatively oriented score, i.e., smaller values indicate better forecasts, and the second representation in (9) implies that it can be expressed in the same unit as the observation. Further, it can be seen as a generalization of the absolute error to probabilistic forecasts and allows comparisons with point forecasts. For the CL0-distribution defined by (1), the CRPS can be computed in closed form (for the exact formula see, Jordan et al., 2019), which enables computationally efficient parameter estimation. For a given lead time,

the forecast skill of competing probabilistic forecasts is compared based on the mean CRPS over all corresponding forecast cases in the verification data.

For assessing the predictive performance of the different forecasts with respect to the binary event that observation x exceeds a given threshold z , one can consider the Brier score (BS; Section 9.4.2 Wilks, 2019), which for a predictive CDF F is defined as

$$\text{BS}(F, x; z) := (F(z) - 1_{\{z \geq x\}})^2 \quad (10)$$

(see e.g., Gneiting and Ranjan, 2011). To compute the Brier Score, the continuous observation x is thus converted to a binary threshold exceedance $1_{\{z \geq x\}}$ with corresponding predicted probability $F(z)$. Note that the Brier score is also negatively oriented and the CRPS is the integral of the BS over all possible thresholds $z \in \mathbb{R}$. For ensemble forecasts, both scores are defined based on the empirical CDF.

For a given probabilistic forecast F , the improvement in a score \mathcal{S}_F with respect to a reference forecast F_{ref} can be quantified with the help of the corresponding skill score defined as

$$\text{SS}_F := 1 - \frac{\bar{\mathcal{S}}_F}{\bar{\mathcal{S}}_{F_{\text{ref}}}},$$

where $\bar{\mathcal{S}}_F$ and $\bar{\mathcal{S}}_{F_{\text{ref}}}$ denote the mean score values over the verification data corresponding to forecasts F and F_{ref} , respectively. In the case studies, we consider the continuous ranked probability skill score (CRPSS) and the Brier skill score (BSS), and will usually use the raw ensemble forecasts as reference model (except for Section B of the Appendix).

A simple tool of assessing the calibration of ensemble forecasts is the verification rank histogram displaying the histogram of ranks of observations with respect to the corresponding ordered ensemble forecasts (Section 9.7.1 Wilks, 2019). For a calibrated K -member ensemble, the ranks should be uniformly distributed on the set $\{1, 2, \dots, K+1\}$. In the case of continuous forecasts specified by predictive distributions, one can either consider the verification rank histogram of simulated ensembles sampled from the predictive PDF, or investigate the probability integral transform (PIT) histogram (Section 9.5.4 Wilks, 2019), where the PIT is the value of the predictive CDF evaluated at the validating observation. For a calibrated predictive distribution, the PIT follows a uniform distribution on the $[0, 1]$ interval. The PIT histogram can thus be considered a continuous counterpart of the verification rank histogram.

Further, one can also investigate the calibration of probabilistic forecasts with the help of the coverage of $(1-\alpha)100\%$, $\alpha \in (0, 1)$, central prediction intervals. Coverage is defined as the proportion of validating observations located between the lower and upper $\alpha/2$ quantiles of the predictive distribution, which for a calibrated probabilistic forecast should be around $(1-\alpha)100\%$. In order to allow a direct comparison with the raw ensemble, the level α is usually chosen to match the nominal coverage of the ensemble range, which for a K -member ensemble equals $(K-1)/(K+1)100\%$.

Finally, point forecasts such as median and mean of the raw ensemble and of the predictive distribution are evaluated with the help of the mean absolute error (MAE) and root mean square error (RMSE). Note that the former is optimal for the median, whereas the latter for the mean (Gneiting, 2011).

4. Case studies

In the following case studies, the forecast skill of various variants of the CL0-EMOS model introduced in Section 3.2 is evaluated. First, we consider a simple CL0-EMOS variant for the HMS AROME-EPS ensemble forecasts of GHI, then we will investigate the performance of the more complex models for the DWD ICON-EPS ensemble forecasts of direct and diffuse irradiance.

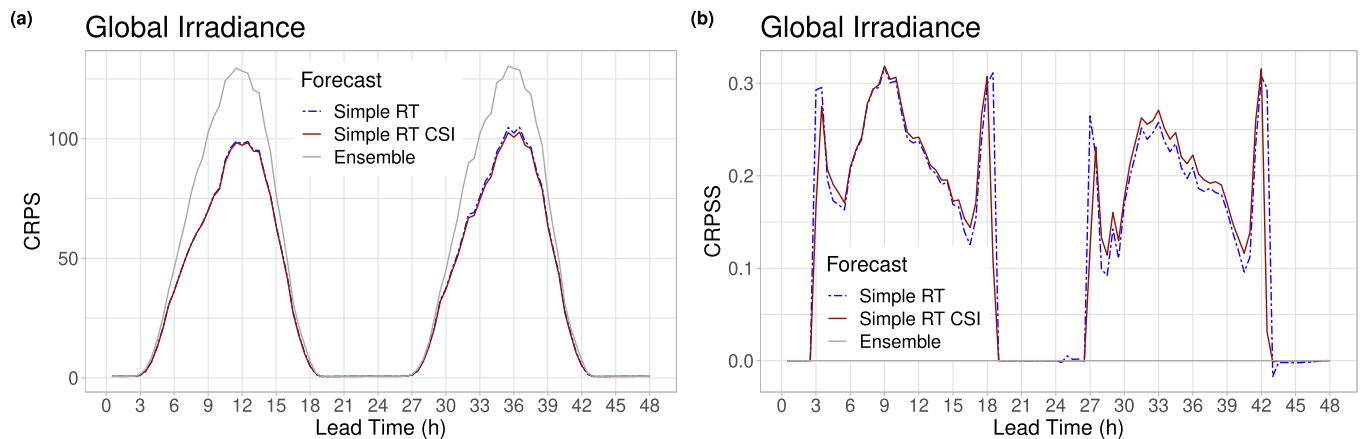


Fig. 2. Mean CRPS of post-processed and raw ensemble forecasts of GHI (a) and CRPSS with respect to the raw ensemble (b) as functions of lead time for the AROME-EPS dataset.

4.1. Results for the AROME-EPS dataset

As discussed in Section 2.1, the AROME-EPS consists of a control member and 10 exchangeable ensemble members obtained using perturbed initial conditions. The dataset at hand covers a short time period only, in particular compared to the ICON-EPS dataset, with forecast-observation pairs available for only 159 calendar days. Therefore, the available training periods cannot be long enough for accurate modeling of seasonal oscillations and we only consider a CLO-EMOS model where the location is linked to the ensemble members via (7) with $K = 2$ and $M_1 = 1$, $M_2 = 10$, which means that six parameters need to be estimated. We consider regional estimation with a rolling training period of

length 31 days, leaving 127 calendar days (9 June 2020 – 13 October 2020) for forecast verification, and refer to this model as the *simple RT* model. The choice of the training period length corresponds to typical values in the post-processing literature and was made to have a similar forecast case per parameter ratio as for the best performing model of Section 4.2. In light of the limited size of the dataset, it is not surprising that the use of monthly expanding training periods or local parameter estimation procedures results in worse predictive performance, and we omit the corresponding results in the interest of brevity. Recall that the ensemble predictions of GHI are provided at a temporal resolution of 30 minutes. As all AROME-EPS forecasts are initialized at 00 UTC, the forecast lead time either coincides with the time of observation or has a

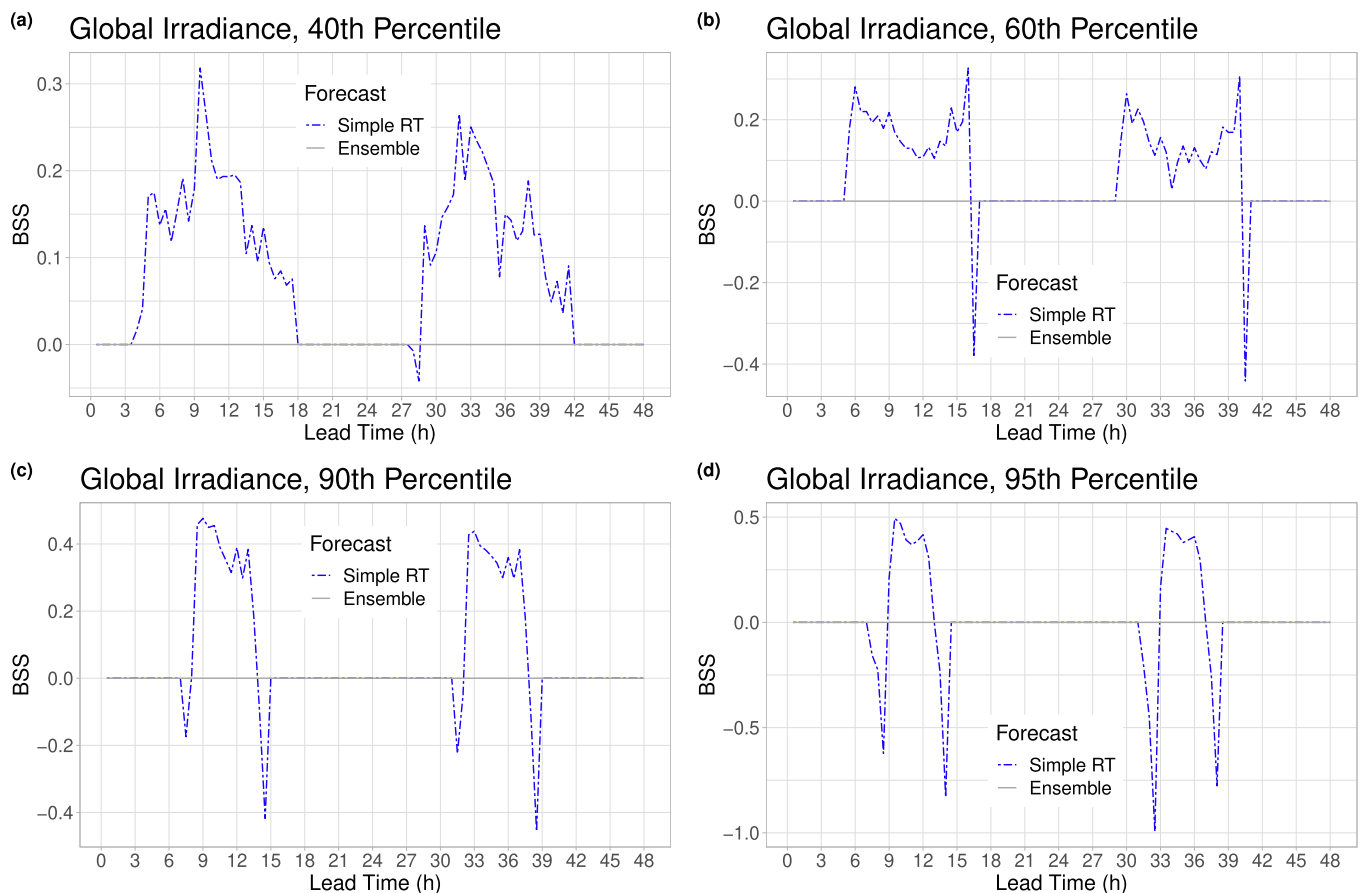


Fig. 3. BSS of post-processed forecasts with respect to the raw ensemble as function of lead time for the AROME-EPS dataset, with thresholds corresponding to the 40th, 60th, 90th and 95th percentiles of observed non-zero GHI.




Fig. 4. Coverage of the nominal 83.33% central prediction intervals of post-processed and raw forecasts (a); MAE of the median forecasts (b) for the AROME-EPS dataset.

shift of 24 hours. Hence, all scores are reported as functions of the lead time. We further average over results from all seven observation locations over Hungary.

Note that in contrast to considering predictions of GHI directly, the standard approach in solar forecasting is the use of a clear-sky index (CSI) as target variable to stationarize the time series of irradiances (Van der Meer et al., 2018; Yang, 2020a; Yang et al., 2020). The clear-sky irradiance used for the normalization is obtained from clear-sky models which estimate the amount of solar radiation arriving at the surface under clear-sky (cloud-free) conditions, see Yang (2020a) for an in-depth discussion and comparison of available models. To investigate the differences between post-processing forecasts of GHI and forecasts of CSI, we follow the procedure outlined in (Section II.A Yang (2020d)) to convert the GHI ensemble predictions f_1, f_2, \dots, f_K and the GHI observation y to CSI values. To do so, we obtained clear-sky irradiance values from the McClellan model using the API provided by the R package `camsRad` Lundström (2016) for the locations and relevant time instances, and converted GHI to CSI by division by the corresponding clear-sky irradiance. We then used an identical model formulation and training procedure as for GHI, and derived 100 equidistant quantiles from the post-processed forecast distributions for CSI. Those quantiles were transformed back to GHI values by multiplying with the corresponding clear-sky irradiance and used for approximating the

verification scores. We refer to this approach as *simple RT CSI*.

Fig. 2 shows the mean CRPS of calibrated and raw ensemble forecasts and the CRPSS with respect to the raw ensemble. Post-processing using the simple RT approach improves the forecast performance when positive irradiance is likely to be observed (3–19 UTC), and performs no corrections otherwise, resulting in a skill score of zero. Note that compared with direct calibration of the GHI, post-processing of the CSI predictions does not result in a substantial difference or clear improvement in forecast skill. These observations are in line with the results reported in Yang (2020d) in a related context. Hence, in the remainder only the results for the former approach will be reported. To assess the statistical significance of the improvements in predictive performance compared to the raw ensemble predictions, we performed a block bootstrap resampling to compute 95% confidence intervals and found that the observed improvements are statistically significant, see Section A.1 of the Appendix. The large jumps in the CRPSS at 4, 19, 27 and 42h are mainly caused by numerical issues as at these lead times the mean CRPS of both raw and post-processed forecasts is very close to 0, and also leads to an increased width of the confidence intervals computed in Section A.1. For qualitatively similar observations in a related context, see e.g. Figure 7 of Bakker et al. (2019). An assessment of the improvements in comparison to a climatological reference forecast is provided in Section B of the Appendix.




Fig. 5. PIT histograms of post-processed and verification rank histograms of raw ensemble forecasts of GHI for the lead times 0–12 h, 12–24 h, 24–36 h and 36–48 h.

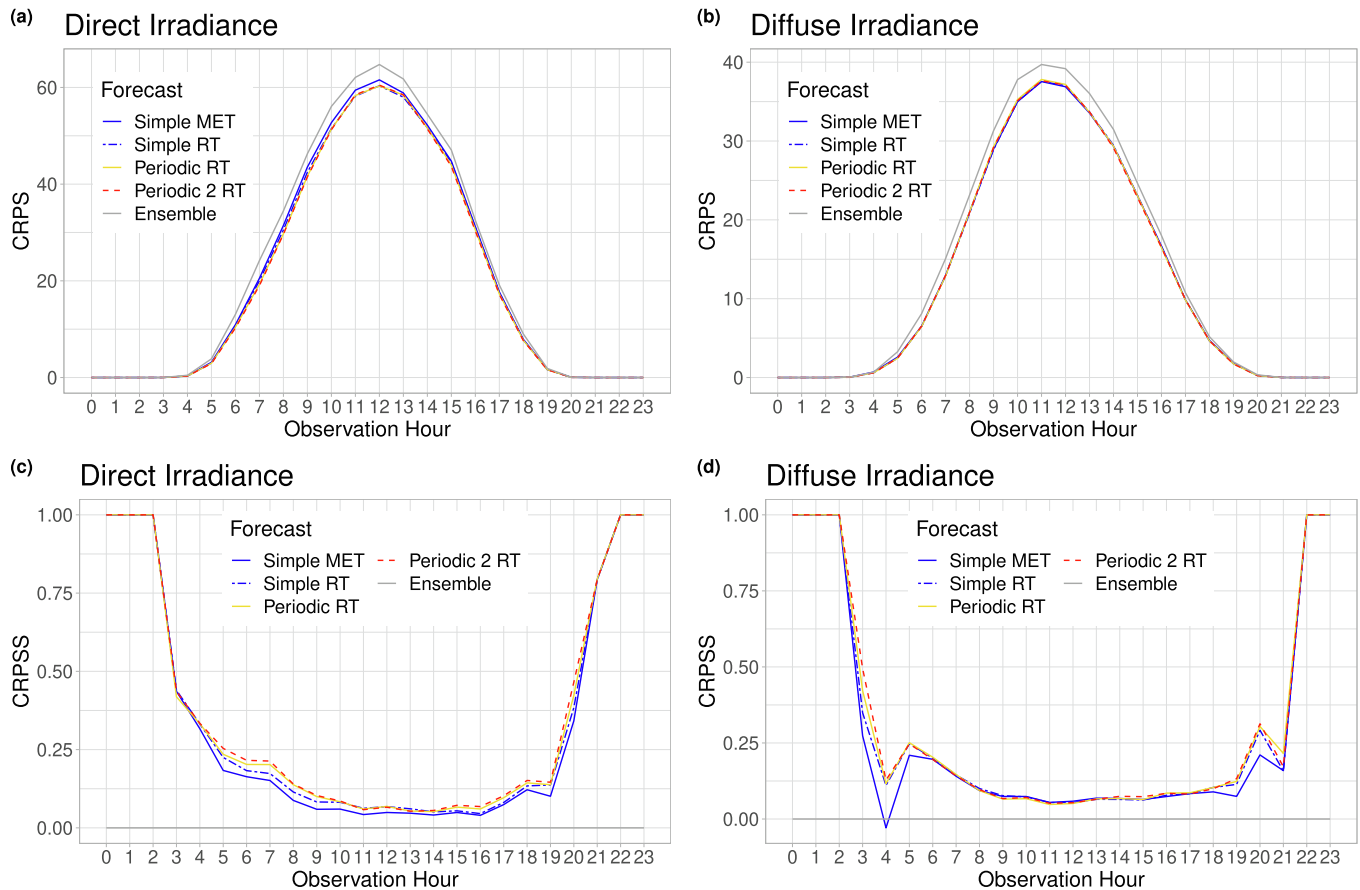


Fig. 6. Mean CRPS of post-processed and raw ensemble forecasts of direct (a) and diffuse (b) irradiance, and corresponding skill scores (c,d) with respect to the raw ensemble as functions of the observation hour for the ICON-EPS dataset.

A similar behavior can be observed for the BSS values shown in Fig. 3, where the threshold values for z correspond to the 40th, 60th, 90th and 95th percentiles of observed non-zero GHI (25, 127, 498, 604 W/m^2). The results are consistent in that the higher the threshold, the shorter the period with a positive mean BS, as the higher thresholds are mostly observed around midday, when the irradiance is strongest. For the corresponding lead times, the post-processed forecasts outperform the raw ensemble. Again, negative skill scores appear only at the boundaries where the mean score values to be compared are very small.

Fig. 4a showing the coverage of the nominal 83.33% central prediction intervals further confirms the improved calibration of the post-processed forecast. Between 3 and 19 UTC, when positive GHI is likely to be observed, the EMOS model results in a coverage close to the nominal value, whereas the coverage of the raw ensemble is consistently below 60%.

Further, Fig. 4b showing the MAE of the median forecasts indicates that post-processing substantially improves the accuracy of point forecasts as well. At the hours of peak irradiance the difference in MAE exceeds 20 W/m^2 . As we will see below, this is in a strong contrast with the results of the second case study (see Fig. 12) and indicates the presence of a bias in the AROME-EPS that is alleviated by post-processing. Similar conclusions can be drawn from the RMSE of the mean forecasts (not shown).

The presence of a bias in the raw ensemble forecasts can also be observed in the verification rank histograms shown in Fig. 5. In addition, the clearly U-shaped verification rank histograms indicate a strong underdispersion, which is in line with the low coverage of the raw ensemble forecasts observed in Fig. 4a and persists across all considered ranges of lead times. However, the ensemble members are more likely to underestimate the true irradiance, which further indicates a negative

bias. Both deficiencies are successfully corrected by statistical post-processing. The PIT histograms of EMOS predictive distributions given in the upper row of Fig. 5 are almost flat indicating just a minor bias for observations between 12:30 and 24 UTC.

4.2. Results for the ICON-EPS dataset

In contrast to the first case study, the ICON-EPS dataset covers a substantially longer time period and therefore allows for considering and comparing more complex model formulations and estimation procedures. Recall that we here consider forecasts of direct irradiance (BHI) and diffuse irradiance (DHI) at temporal resolutions of 1 h (for lead times up to 48 h), 3 h (for lead times 51–72 h) and 6 h (for lead times 78–120 h), resulting in a forecast horizon of 120 h.

As members of the ICON-EPS are obtained with the help of random perturbations, they can be regarded as exchangeable. Hence, for post-processing we use the CLO-EMOS model with locations linked to the ensemble members either via (7) or via (8) with $K = 1$. Thus, for model (3) with location (7) (which was the only model variant considered in Section 4.1 and is referred to as *simple model*) one has to estimate five unknown parameters, whereas more complex approaches, which account for seasonal variations in the link function (8) of the location parameter via (4) (referred to as *periodic model*) or (5) (referred to as *periodic 2 model*), require the estimation of a total of 11 and 15 parameters, respectively.

The period from 27 December 2018 - 31 December 2019 is used for training purposes only, the calendar year 2020 (366 calendar days) for model verification, which leaves enough flexibility for choosing a sufficiently long training period even for local modeling. Two different training configurations are investigated: a rolling training (RT) period of

length 365 days, and a monthly expanding training (MET) scheme, where all data until the end of the last month before the forecast date under consideration is used for training. In the latter case, the first training period includes all data prior to calendar year 2020. According to initial studies (not shown), MET provides reasonable verification scores only for the simple model. Therefore, we report results for the simple model with rolling (*simple RT*) and monthly expanding training (*simple MET*) as well as for the periodic models with rolling training (*periodic RT* and *periodic 2 RT*). Given the negligible differences we observed when comparing post-processing of GHI and CSI for the AROME-EPS data in Section 4.1, we only consider predictions of BHI and DHI without normalization by the corresponding clear-sky irradiances here.

Raw ensemble forecasts of direct and diffuse irradiance are used as references models. Unless indicated otherwise, results discussed below are averaged over all three observation locations and all four initialization times of the NWP model. Note that this might make the interpretation of the results more involved than in the first case study due to the interacting effects of forecast initialization time, lead time, and corresponding time of day of the observation.

First, we investigate diurnal effects by examining the dependence of the mean CRPS of the various forecast models on the time of the observation shown in Fig. 6a,b. In order to provide a fair comparison, we take only the first 48h of the forecast horizon into account, where hourly forecasts are available. Both for BHI and DHI, all post-processing

methods outperform the raw ensemble forecasts at all time points when positive irradiance is likely to be observed. According to the skill scores with respect to the raw ensemble shown in Fig. 6c, in the case of direct irradiance the predictive performance mainly depends on the complexity of model formulations and parameter estimation, with more complex models exhibiting better forecast performance. However, the differences between the various EMOS approaches are relatively minor. The same applies for diffuse irradiance in early and late hours (see Fig. 6d), whereas between 6 and 18 UTC there is no visible difference in the skill of the different EMOS models. For corresponding results against a climatological reference model, see Section B of the Appendix.

Note that the apparent periodic oscillations in the CRPS values might be partly caused by the pooling of different observation hours due to the four considered initialization times. In contrast to the AROME-EPS, post-processing also improves the predictive performance at night achieving a CRPS of almost zero. ICON-EPS fails to achieve mean CRPS values of zero due to occasional predictions of non-zero irradiance values during night times.

Fig. 7a, which shows the CRPS with respect to the raw BHI ensemble forecasts as function of the lead time, confirms the observations from Fig. 6c. Here the differences are more pronounced due to the different scaling of the vertical axis, and again, the periodic 2 model with rolling training period exhibits the best forecast skill, whereas the simple model with monthly expanding training shows the smallest CRPS. In general, all skill scores decrease for longer lead times, which

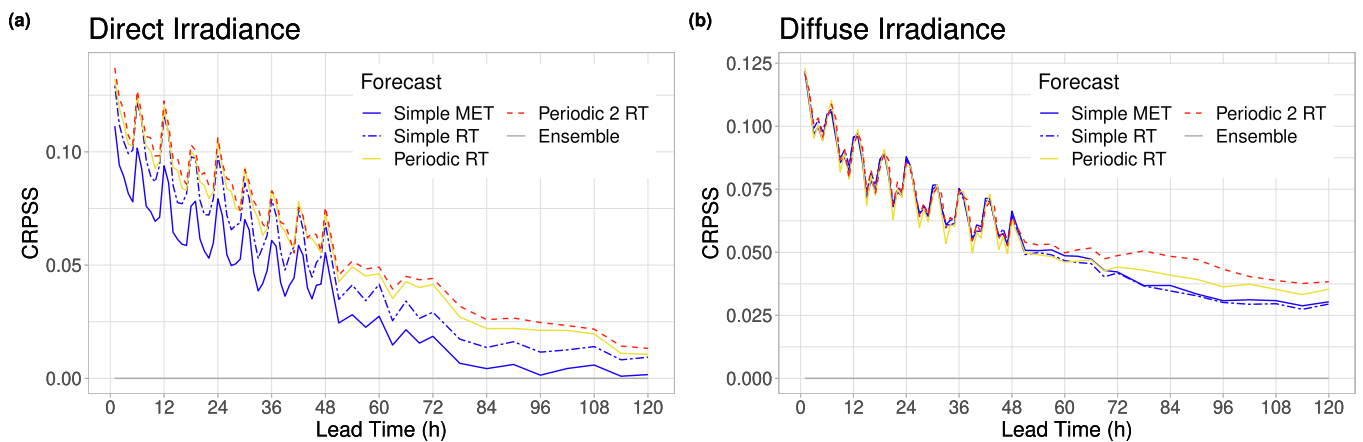


Fig. 7. CRPS of post-processed forecasts of direct (a) and diffuse (b) irradiance with respect to the raw ensemble as functions of the lead time for the ICON-EPS dataset.

Table 1

Overall CRPS and CRPS for individual locations of post-processed forecasts of direct and diffuse irradiance with respect to the raw ensemble.

Lead Time	Model	Direct Irradiance				Diffuse Irradiance			
		Overall	Karlsruhe	Berlin	Hamburg	Overall	Karlsruhe	Berlin	Hamburg
1–24 h	Simple MET	0.076	0.114	0.029	0.082	0.089	0.099	0.075	0.093
	Simple RT	0.095	0.130	0.070	0.083	0.091	0.100	0.080	0.092
	Periodic RT	0.101	0.133	0.077	0.090	0.089	0.099	0.080	0.086
	Periodic 2 RT	0.104	0.132	0.082	0.097	0.091	0.097	0.083	0.092
25–48 h	Simple MET	0.050	0.091	-0.000	0.054	0.066	0.076	0.050	0.070
	Simple RT	0.064	0.102	0.032	0.054	0.066	0.075	0.053	0.069
	Periodic RT	0.072	0.102	0.045	0.066	0.064	0.071	0.055	0.063
	Periodic 2 RT	0.074	0.099	0.047	0.073	0.066	0.069	0.060	0.069
51–72 h	Simple MET	0.021	0.062	-0.018	0.018	0.048	0.055	0.041	0.046
	Simple RT	0.033	0.071	0.008	0.020	0.046	0.053	0.041	0.043
	Periodic RT	0.043	0.075	0.021	0.030	0.047	0.053	0.045	0.041
	Periodic 2 RT	0.046	0.069	0.024	0.043	0.051	0.049	0.055	0.049
73–120 h	Simple MET	0.004	0.028	-0.016	-0.002	0.032	0.032	0.039	0.025
	Simple RT	0.013	0.036	-0.001	0.001	0.031	0.031	0.040	0.021
	Periodic RT	0.019	0.032	0.017	0.008	0.038	0.037	0.048	0.026
	Periodic 2 RT	0.022	0.024	0.021	0.022	0.043	0.036	0.062	0.030

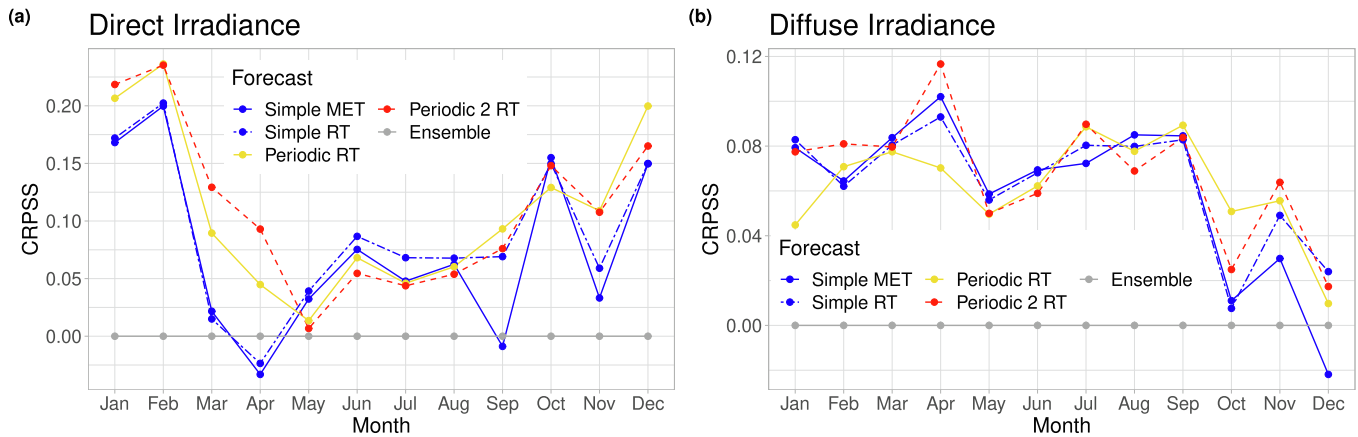


Fig. 8. CRPSS of post-processed forecasts of direct (a) and diffuse (b) irradiance with respect to the raw ensemble, computed based on monthly mean values for the ICON-EPS dataset.

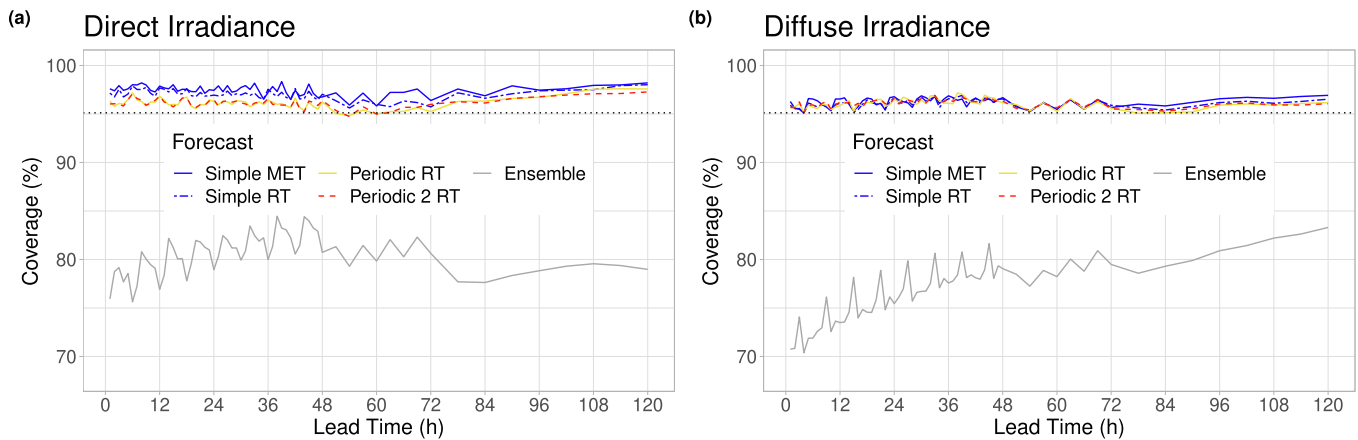


Fig. 9. Coverage of nominal 95.12% central prediction intervals of post-processed and raw forecasts of direct (a) and diffuse (b) irradiance for the ICON-EPS dataset.

also holds for the corresponding CRPSS values for DHI (Fig. 7b). Overall, slightly larger improvements relative to the raw ensemble are observed for direct than for diffuse irradiance, and none of the models result in negative skill scores. Up to a lead time of 48h, there are no visible differences between the various EMOS approaches. For longer lead times, similar to BHI, the most complex periodic 2 model shows the best predictive performance, whereas the simple model with parameters estimated using a rolling training period is now the least skillful. Recall that for longer lead times the forecasts refer to a longer time period and thus seasonal effects regarding the diurnal cycle might be captured by the more complex models.

An assessment of the statistical significance of the improvements in predictive performance compared to the raw ensemble predictions shows that in the case of BHI, post-processing results in a significant improvement in mean CRPS up to 60h ahead, whereas post-processed forecasts of DHI significantly outperform the raw ensemble over the entire forecast horizon of 120h. The corresponding results and more details can be found in Section A.2 of the Appendix.

A third aspect is the dependence of the forecast skill on the observation location. Table 1 shows the overall CRPSS of the different EMOS models with respect to the raw forecasts and the corresponding CRPSS values of the three different cities for four different intervals of the forecast horizon. The main message of these results is that the magnitude of improvements in predictive performance resulting from post-processing strongly depends on the location. For both variables, Karlsruhe benefits the most, while for Berlin after 24h, and for Hamburg after

78h the simple MET model performs worse than the raw BHI ensemble forecast and results in negative skill scores. Among the competing models for BHI the most complex periodic 2 RT model shows the best forecast skill for Berlin and Hamburg, and shows the best overall performance as well. In the case of DHI, the differences in performance between the various EMOS models are much smaller, which is in line with the results observed in Fig. 7b. In particular, none of the more complex models consistently outperforms the simple MET model.

To investigate seasonal effects in the improvements achieved via post-processing, Fig. 8 shows the CRPSS of the post-processed forecasts based on monthly mean values. For direct irradiance, the improvements are generally larger in winter than in summer. From November to April, the differences among the post-processing approaches are most pronounced, and more complex model formulations that incorporate seasonal effects particularly show improved performance. For diffuse irradiance, the overall level of improvements in terms of the mean CRPS is smaller (note the different scale of the vertical axes). Only minor seasonal effects in the form of smaller improvements between October and December can be detected.

To simplify the presentation of the results, in the remaining part of this section we consider pooled data of all locations, months and observation hours and display the dependence on the lead time only.

The improved calibration of post-processed forecasts can also be observed in the coverage plots of Fig. 9. All post-processing approaches result in a coverage close to the nominal 95.12% for all lead times, whereas the maximal coverage of the raw ensemble is below 85% for both

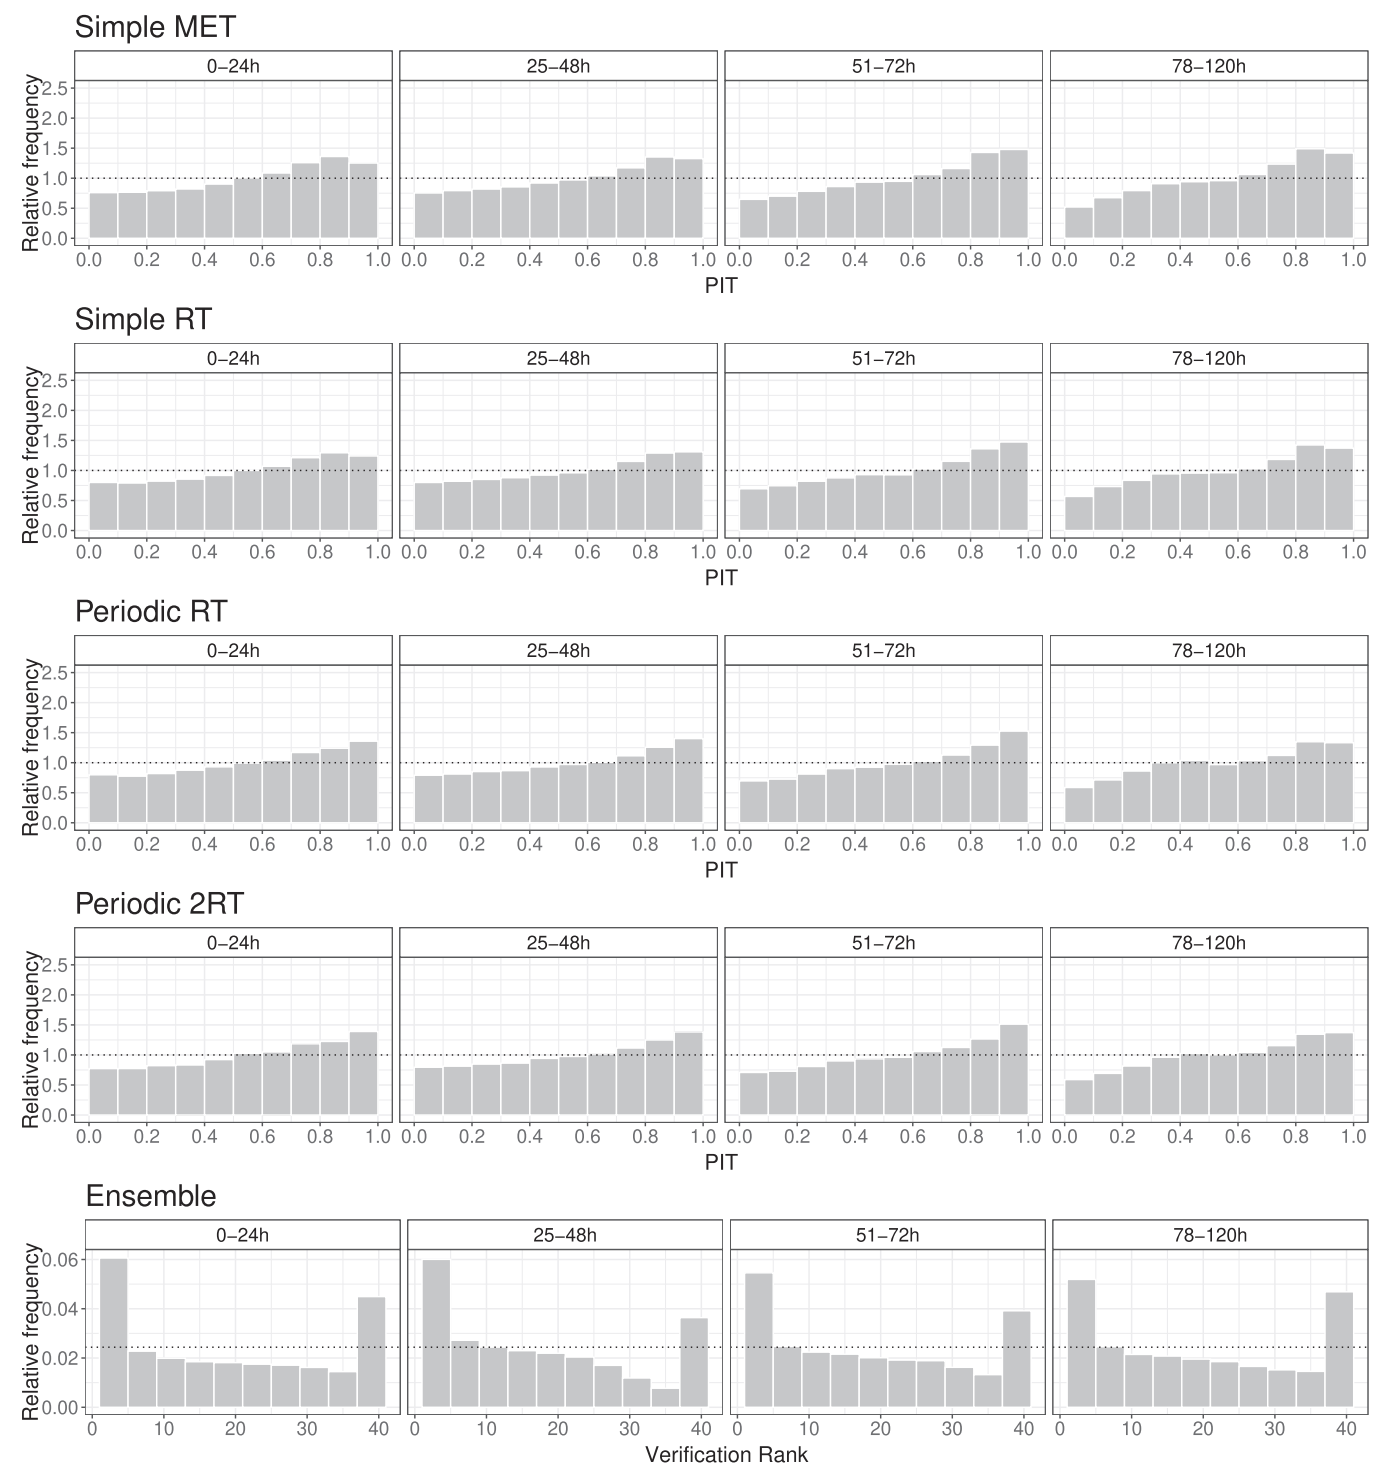


Fig. 10. PIT histograms of post-processed and verification rank histograms of raw ensemble forecasts of DNI for the lead times 0–24 h, 25–48 h, 51–72 h and 78–120 h for the ICON-EPS dataset.




Fig. 11. PIT histograms of post-processed and verification rank histograms of raw ensemble forecasts of DHI for the lead times 0–24 h, 25–48 h, 51–72 h and 78–120 h for the ICON-EPS dataset.

variables. The difference between post-processed forecasts of BHI is more pronounced with periodic models being the closest to the nominal value. These results are in line with the shapes of the PIT and verification rank histograms of Figs. 10 and 11. Raw ensemble forecasts of BHI are strongly underdispersive and slightly biased for all lead times. Even though this is slightly alleviated for longer lead times, the PIT histograms of all EMOS

models are much closer to the desired uniform distribution. However, some bias still remains in the post-processed forecasts. In contrast, neither the PIT histograms of post-processed, nor the verification rank histograms of raw forecasts of DHI indicate any bias (Fig. 11), and all EMOS approaches successfully correct the underdispersion of the raw ensemble resulting in almost perfectly uniform PIT histograms.




Fig. 12. MAE of the median forecasts of direct (a) and diffuse (b) irradiance.

Finally, Fig. 12 showing the MAE of the median forecasts indicates that while post-processing substantially improves the calibration of probabilistic forecasts, it has a minor effect on the accuracy of point forecasts. The difference in MAE is less than 2 W/m^2 for direct irradiance and 0.6 W/m^2 for diffuse irradiance for all considered lead times. The sharp changes in MAE values at 51h and 78h are results of the change in temporal resolution of the forecasts. Corresponding results for the RMSE of the mean forecasts are very similar and thus not shown here.

5. Conclusions

We propose a post-processing method for ensemble weather predictions of solar irradiance where probabilistic forecasts are obtained in the form of a logistic distribution left-censored at zero. Several model variants that differ in terms of the temporal composition of training datasets and adjustments to seasonal variations in the model formulation are evaluated in two case studies. Even though the case studies cover distinct geographical regions, NWP systems, types of solar irradiance and temporal resolutions, the results presented in Section 4 indicate that the proposed post-processing models are able to consistently and significantly improve the forecast performance of the raw ensemble predictions up to lead times of at least 48 hours. The improvements from post-processing are larger for the AROME-EPS dataset, possibly due to a lower skill of the raw ensemble predictions resulting from a bias in addition to the observed underdispersion. For the ICON-EPS dataset, we observed that more complex post-processing models tend to show better predictive performances, but the differences between model variations rarely show a high level of statistical significance. For the GHI predictions of the AROME-EPS dataset, we only found negligible differences when comparing post-processing models for GHI and CSI. This is in line with the results reported in Yang (2020d) and suggests that the standard practice of normalizing the irradiance forecasts by clear-sky irradiance does not lead to improvements in forecast performance here.

The overall level of improvements achieved via statistical post-processing of the solar irradiance forecasts of the raw ensemble are comparable to meteorological variables such as precipitation accumulation (Scheuerer, 2014; Baran and Nemoda, 2016) or total cloud cover (Baran et al., 2021) in case of the ICON-EPS dataset, and slightly larger for the AROME-EPS data. Post-processing ensemble predictions of those variables is often seen as a more difficult task compared to variables

such as temperature (Gneiting et al., 2005) or wind speed (Thorarinsdottir and Gneiting, 2010) where substantially larger improvements can be achieved. Nonetheless, the observed improvements are statistically significant for lead times of up to 2 days, and will likely be of relevance for solar energy forecasting in terms of potential economic benefits and improved balancing of demand and supply for integrating volatile PV power systems into the electrical grid. The datasets used in the two case studies are somewhat limited in terms of their temporal extent, in particular the AROME-EPS data. An interesting aspect for future work will be to for example compare different ways of accounting for seasonal variability once longer, ideally multi-year periods of data have become available.

The post-processing models for solar irradiance considered here provide several avenues for future work. Regarding the general model setup, we have only used ensemble predictions of the target variable as inputs to the post-processing model. However, a recent focus of the post-processing literature has been the use of modern machine learning methods that allow one to incorporate arbitrary predictor variables and to model possibly nonlinear relations to the forecast distribution parameters (see for a recent review, Vannitsem et al., 2021). The EMOS models considered here could for example be extended following the gradient boosting extension framework proposed in Messner et al. (2017) or the neural network-based approach of Rasp and Lerch (2018). For similar considerations in the solar irradiance forecasting literature where additional predictors from NWP model output are used, albeit for different types for probabilistic forecasting methods, see, for example, Sperati et al. (2016) and Bakker et al. (2019).

Further, we have restricted our attention to univariate forecasts for a single location, lead time and target variable. However, in particular in the context of energy forecasting, many practical applications require an accurate modeling of spatial, temporal, or inter-variable dependencies (Pinson and Messner, 2018). A large variety of multivariate post-processing methods has been proposed over the past years (see for a recent overview, Lerch et al., 2020), and a comparison of those approaches in the context of solar energy forecasting might be an interesting starting point for future research.

Finally, the development of post-processing models for solar irradiance was motivated by the aim of improving probabilistic solar energy forecasting. To that end, it would be interesting to investigate the effect of post-processing NWP ensemble forecasts of solar irradiance for PV power prediction, and for example compare to direct probabilistic

models of PV power output (see e.g. Alessandrini et al., 2015). In a related study in the context of wind energy, Phipps et al. (2020) find that a two-step strategy of post-processing both wind and power ensemble forecasts performs best and that the calibration of the power predictions constitutes a crucial step. Ideally, statistical post-processing of solar irradiance forecasts could contribute an important component to modern, fully integrated renewable energy forecasting systems (see e.g., Haupt et al., 2020).

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Significance of improvement in forecast performance

To get an insight into the uncertainty and statistical significance of the verification scores, we accompany the CRPSS values of the best performing EMOS models (simple RT for the AROME-EPS dataset and periodic 2 RT for the ICON-EPS dataset) with 95% confidence intervals. The corresponding standard deviations are obtained from 2000 block bootstrap samples calculated using the stationary bootstrap scheme, where the mean block length is computed according to Politis and Romano (1994).

A.1. CRPSS of the AROME-EPS

In the case of the AROME-EPS dataset we consider only the simple EMOS model with a 31 day rolling training period. Fig. 13 is an enhanced version of Fig. 2b where corresponding 95% confidence intervals for the CRPSS are added. The skill scores of post-processed forecasts are significantly positive for all time periods where positive irradiance is likely to be observed.




Fig. 13. CRPSS of EMOS post-processed forecasts with respect to the raw ensemble together with 95% confidence intervals for the AROME-EPS dataset.

A.2. CRPSS of the ICON-EPS

Fig. 14 shows the CRPSS of the best performing periodic 2 RT EMOS model with respect to the raw ICON-EPS forecast and the corresponding 95% confidence intervals. In the case of direct irradiance the improvement in mean CRPS is significant up to 60h, whereas for diffuse irradiance it is significant for all considered lead times. Further, comparing Figs. 14 and 7 it can be observed (especially in the case of diffuse irradiance), that in terms of mean CRPS there is no significant difference between the various post-processing methods.

Acknowledgments

Benedikt Schulz and Sebastian Lerch gratefully acknowledge support by the Deutsche Forschungsgemeinschaft through SFB/TRR 165 “Waves to Weather”. Sándor Baran was supported by the National Research, Development and Innovation Office under Grant No. NN125679. He also acknowledges the support of the EFOP-3.6.2-16-2017-00015 project, which was co-financed by the Hungarian Government and the European Social Fund. The authors thank Gabriella Szépszó and Mihály Szűcs from the HMS for providing the AROME-EPS data. We thank DWD for providing access to ICON-EPS and station observation data via the Open Data Server of DWD (<https://opendata.dwd.de/>), and Marco Wurth for assistance in obtaining the forecast data. Constructive comments on an earlier version of the manuscript by two anonymous referees are gratefully acknowledged.

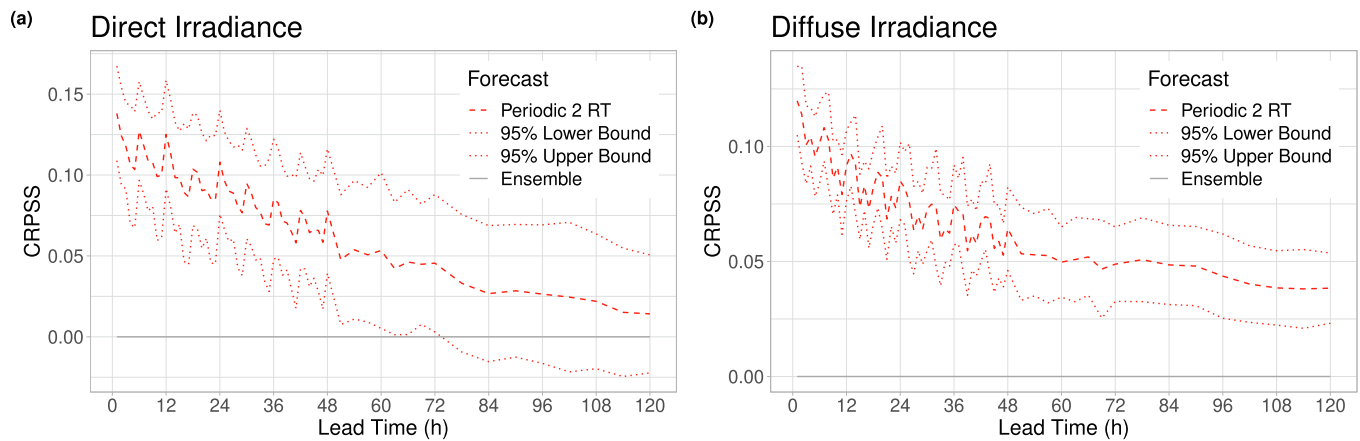


Fig. 14. CRPSS of the best performing post-processed forecasts of direct (a) and (b) diffuse irradiance with respect to the raw ensemble together with 95% confidence intervals for the ICON-EPS dataset.

Appendix B. Predictive performance with respect to climatology

The relative improvements of the post-processed forecasts in terms of the CRPSS were investigated in Section 4 by computing the corresponding skill score values of the CRPSS using the ensemble forecast as reference (see Figs. 2 and 6c,d). To further investigate the relative improvements in comparison to a more naïve reference, Figs. 15 and 16 show corresponding values of the CRPSS using climatological forecast as reference model, where observations of the rolling training period are considered as a forecast ensemble. This can be viewed as a persistence ensemble in the terminology of Yang (2019).




Fig. 15. CRPSS of EMOS post-processed and raw ensemble forecasts with respect to climatology as functions of the observation hour for the AROME-EPS dataset.

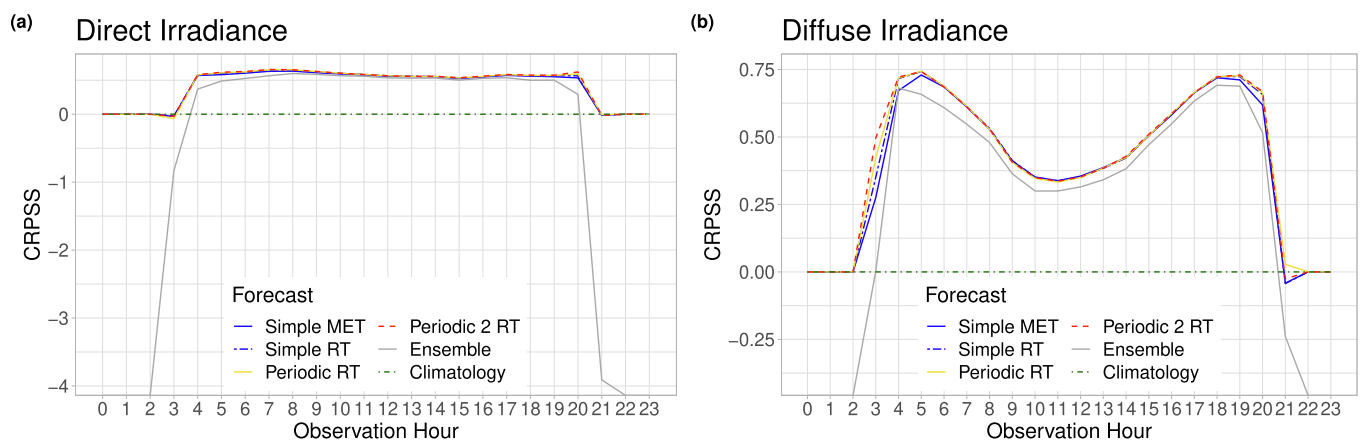


Fig. 16. CRPSS of post-processed and raw ensemble forecasts of direct (a) and (b) diffuse irradiance with respect to climatology as functions of the observation hour for the ICON-EPS dataset.

In the case of the AROME-EPS individual climatological forecasts are based on observations of the preceding 31 days, whereas for the ICON-EPS a climatological ensemble has 365 members. Both ensemble and the post-processed forecasts show clear improvements for times of day during which it is unlikely to observe zero irradiance. Corresponding skill scores for MAE and RMSE indicate a very similar behavior of the deterministic forecasts and are not shown here in the interest of brevity.

References

- Alessandrini, S., Delle Monache, L., Sperati, S., Cervone, G., 2015. An analog ensemble for short-term probabilistic solar power forecast. *Appl. Energy* 157, 95–110.
- Bakker, K., Whan, K., Knap, W., Schmeits, M., 2019. Comparison of statistical post-processing methods for probabilistic NWP forecasts of solar radiation. *Sol. Energy* 191, 138–150.
- Baran, S., Lerch, S., 2016. Mixture EMOS model for calibrating ensemble forecasts of wind speed. *Environmetrics* 27, 116–130.
- Baran, S., Lerch, S., 2018. Combining predictive distributions for the statistical post-processing of ensemble forecasts. *Int. J. Forecast.* 34, 477–496.
- Baran, A., Lerch, S., El Ayari, M., Baran, S., 2021. Machine learning for total cloud cover prediction. *Neural Comput. Appl.* 33, 2605–2620.
- Baran, S., Nemoda, D., 2016. Censored and shifted gamma distribution based EMOS model for probabilistic quantitative precipitation forecasting. *Environmetrics* 27, 280–292.
- Bauer, P., Thorpe, A., Brunet, G., 2015. The quiet revolution of numerical weather prediction. *Nature* 525, 47–55.
- Becker, R., Behrens, K., 2012. Quality assessment of heterogeneous surface radiation network data. *Adv. Sci. Res.* 8, 93–97.
- Bentzien, S., Friederichs, P., 2012. Generating and calibrating probabilistic quantitative precipitation forecasts from the high-resolution NWP model COSMO-DE. *Weath. Forecast.* 27, 988–1002.
- DWD Climate Data Center, 2020. Recent 10-minute station observations of solar incoming radiation, longwave downward radiation and sunshine duration for Germany. last accessed: January 15, 2021.
- Fraley, C., Raftery, A.E., Gneiting, T., 2010. Calibrating multimodel forecast ensembles with exchangeable and missing members using bayesian model averaging. *Mon. Weather Rev.* 138, 190–202.
- Fraunhofer Institute for Solar Energy Systems, 2020. Recent Facts about Photovoltaics in Germany. Technical report, available at <https://www.ise.fraunhofer.de/en/publications/studies/recent-facts-about-pv-in-germany.html>.
- Gneiting, T., 2011. Making and evaluating point forecasts. *J. Am. Stat. Assoc.* 106, 746–762.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* 102, 359–378.
- Gneiting, T., Raftery, A.E., Westveld III, A.H., Goldman, T., 2005. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Weather Rev.* 133, 1098–1118.
- Gneiting, T., Ranjan, R., 2011. Comparing density forecasts using threshold- and quantile-weighted scoring rules. *J. Bus. Econ. Stat.* 29, 411–422.
- González Ordiano, J.A., Gröll, L., Mikut, R., Hagenmeyer, V., 2020. Probabilistic energy forecasting using the nearest neighbors quantile filter and quantile regression. *Int. J. Forecast.* 36, 310–323.
- Gottwalt, S., Gärtner, J., Schmeck, H., Weinhardt, C., 2016. Modeling and valuation of residential demand flexibility for renewable energy integration. *IEEE Trans. Smart Grid* 8, 2565–2574.
- Haupt, S.E., Casado, M.G., Davidson, M., Dobschinski, J., Du, P., Lange, M., Miller, T., Mohren, C., Motley, A., Pestana, R., 2019. The use of probabilistic forecasts: Applying them in theory and practice. *IEEE Power Energy Mag.* 17, 46–57.
- Haupt, S.E., McCandless, T.C., Dettling, S., Alessandrini, S., Lee, J.A., Linden, S., Petzke, W., Brummet, T., Nguyen, N., Kosović, B., 2020. Combining artificial intelligence with physics-based methods for probabilistic renewable energy forecasting. *Energies* 13, 1979.
- Hemri, S., Scheuerer, M., Pappenberger, F., Bogner, K., Haiden, T., 2014. Trends in the predictive performance of raw ensemble weather forecasts. *Geophys. Res. Lett.* 41, 9197–9205.
- Hong, T., Fan, S., 2016. Probabilistic electric load forecasting: A tutorial review. *Int. J. Forecast.* 32 (3), 914–938.
- Hong, T., Pinson, P., Wang, Y., Weron, R., Yang, D., Zareipour, H., 2020. Energy forecasting: A review and outlook. *IEEE Open Access J. Power Energy* 7, 376–388.
- Jávorné Radnóczi, K., Várkonyi, A., Szépszó, G., 2020. On the way towards the AROME nowcasting system in Hungary. *ALADIN-HIRLAM Newsletter* 14, 65–69.
- Jordan, A., Krüger, F., Lerch, S., 2019. Evaluating probabilistic forecasts with scoringRules. *J. Stat. Softw.* 90, 1–37.
- La Salle, J.L.G., Badosa, J., David, M., Pinson, P., Lauret, P., 2020. Added-value of ensemble prediction system on the quality of solar irradiance probabilistic forecasts. *Renew. Energy* 162, 1321–1339.
- Lang, M.N., Lerch, S., Mayr, G.J., Simon, T., Stauffer, R., Zeileis, A., 2020. Remember the past: a comparison of time-adaptive training schemes for non-homogeneous regression. *Nonlinear Process. Geophys.* 27, 23–34.
- Lauret, P., David, M., Pinson, P., 2019. Verification of solar irradiance probabilistic forecasts. *Sol. Energy* 194, 254–271.
- Lerch, S., Baran, S., Möller, A., Groß, J., Schefzik, R., Hemri, S., Graeter, M., 2020. Simulation-based comparison of multivariate ensemble post-processing methods. *Nonlinear Process. Geophys.* 27, 349–371.
- Lundstrom, L., 2016. camsRad: Client for CAMS Radiation Service. R package version 0.3.0. <https://CRAN.R-project.org/package=camsRad>.
- Messner, J.W., Mayr, G.J., Zeileis, A., 2017. Nonhomogeneous boosting for predictor selection in ensemble postprocessing. *Mon. Weather Rev.* 145, 137–147.
- Phipps, K., Lerch, S., Andersson, M., Mikut, R., Hagenmeyer, V., Ludwig, N., 2020. Evaluating ensemble post-processing for wind power forecasts. Preprint, available at <http://arxiv.org/abs/2009.14127>.
- Pinson, P., Messner, J.W., 2018. Application of postprocessing for renewable energy. In: Vannitsem, S., Wilks, D.S., Messner, J.W. (Eds.), *Statistical Postprocessing of Ensemble Forecasts*. Elsevier, pp. 241–266.
- Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. *J. Am. Stat. Assoc.* 89, 1303–1313.
- Rasp, S., Lerch, S., 2018. Neural networks for postprocessing ensemble weather forecasts. *Mon. Weather Rev.* 146 (11), 3885–3900.
- Reinert, D., Prill, F., Frank, H., Denhard, M., Baldauf, M., Schraff, C., Gebhardt, C., Marsigli, C., Zängl, G., 2020. DWD Database Reference for the Global and Regional ICON and ICON-EPS Forecasting System. Version 2.1.1. Deutscher Wetterdienst, Offenbach am Main.
- Scheuerer, M., 2014. Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Q. J. R. Meteorol. Soc.* 140, 1086–1096.
- Sloughter, J.M.L., Raftery, A.E., Gneiting, T., Fraley, C., 2007. Probabilistic quantitative precipitation forecasting using Bayesian model averaging. *Mon. Weather Rev.* 135, 3209–3220.
- Sperati, S., Alessandrini, S., Delle Monache, L., 2016. An application of the ECMWF Ensemble Prediction System for short-term solar power forecasting. *Sol. Energy* 133, 437–450.
- Thorarindottir, T.L., Gneiting, T., 2010. Probabilistic forecasts of wind speed: Ensemble model output statistics by using heteroscedastic censored regression. *J. Roy. Stat. Soc. Ser. A (Stat. Soc.)* 173, 371–388.
- Van der Meer, D.W., Widen, J., Munkhammar, J., 2018. Review on probabilistic forecasting of photovoltaic power production and electricity consumption. *Renew. Sustain. Energy Rev.* 81, 1484–1512.
- 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., Hietala, L., Mestre, O., Moret, L., Odak Plenković, I., Schmeits, M., Taillardat, M., Van den Bergh, J., Van Schaeybroeck, B., Whan, K., Ylhäisi, J., 2021. Statistical postprocessing for weather forecasts – review, challenges and avenues in a big data world. *Bull. Am. Meteorol. Soc.* 102, E681–E699.
- Wilks, D.S., 2018. *Univariate ensemble postprocessing*. In: Vannitsem, S., Wilks, D.S., Messner, J.W. (Eds.), *Statistical Postprocessing of Ensemble Forecasts*. Elsevier, pp. 49–89.
- Wilks, D.S., 2019. *Statistical Methods in the Atmospheric Sciences*, 4th Edition. Elsevier Academic Press.
- Yagli, G.M., Yang, D., Srinivasan, D., 2020. Ensemble solar forecasting using data-driven models with probabilistic post-processing through GAMLSS. *Sol. Energy* 208, 612–622.
- Yang, D., 2019. A guideline to solar forecasting research practice: Reproducible, operational, probabilistic or physically-based, ensemble, and skill (ROPES). *J. Renew. Sustain. Energy* 11, 022701.
- Yang, D., 2020a. Choice of clear-sky model in solar forecasting. *J. Renew. Sustain. Energy* 12, 026101.
- Yang, D., 2020b. Comment: Operational aspects of solar forecasting. *Sol. Energy* 210, 38–40.
- Yang, D., 2020c. Ensemble model output statistics as a probabilistic site-adaptation tool for satellite-derived and reanalysis solar irradiance. *J. Renew. Sustain. Energy* 12, 016102.
- Yang, D., 2020d. Ensemble model output statistics as a probabilistic site-adaptation tool for solar irradiance: A revisit. *J. Renew. Sustain. Energy* 12, 036101.
- Yang, D., Alessandrini, S., Antonanzas, J., Antonanzas-Torres, F., Badescu, V., Beyer, H. G., Blaga, R., Boland, J., Bright, J.M., Coimbra, C.F., 2020. Verification of deterministic solar forecasts. *Sol. Energy* 210, 20–37.
- Yang, D., van der Meer, D., 2021. Post-processing in solar forecasting: Ten overarching thinking tools. *Renew. Sustain. Energy Rev.* 140, 110735.
- Zängl, G., Reinert, D., Rípodas, P., Baldauf, M., 2015. The ICON (ICOsaedral Non-hydrostatic) modelling framework of DWD and MPI-M: Description of the non-hydrostatic dynamical core. *Q. J. R. Meteorol. Soc.* 141, 563–579.
- Zelikman, E., Zhou, S., Irvin, J., Raterink, C., Sheng, H., Kelly, J., Rajagopal, R., Ng, A. Y., Gagne, D., 2020. Short-term solar irradiance forecasting using calibrated probabilistic models. *NeurIPS Workshop on Tackling Climate Change with Machine Learning*, <https://www.climatechange.ai/papers/neurips2020/6>.