

RESEARCH ARTICLE

Enhancing multivariate post-processed visibility predictions utilizing Copernicus Atmosphere Monitoring Service forecasts

Mária Lakatos^{1,2} | Sándor Baran¹ 

¹Faculty of Informatics, University of Debrecen, Debrecen, Hungary

²Doctoral School of Informatics, University of Debrecen, Debrecen, Hungary

Correspondence

Sándor Baran, Faculty of Informatics, University of Debrecen, Kassai út 26, H-4028 Debrecen, Hungary.
Email: baran.sandor@inf.unideb.hu

Funding information

National Research, Development and Innovation Office, Grant/Award Number: K142849; Nemzeti Kutatási, Fejlesztési és Innovációs Alap, Grant/Award Number: ÚNKP-23-3

Abstract

In our contemporary era, meteorological weather forecasts increasingly incorporate ensemble predictions of visibility—a parameter of great importance in aviation, maritime navigation, and air quality assessment, with direct implications for public health. However, this weather variable falls short of the predictive accuracy achieved for other quantities issued by meteorological centers. Therefore, statistical post-processing is recommended to enhance the reliability and accuracy of predictions. By estimating the predictive distributions of the variables with the aid of historical observations and forecasts, one can achieve statistical consistency between true observations and ensemble predictions. Visibility observations, following the recommendation of the World Meteorological Organization, are typically reported in discrete values; hence, the predictive distribution of the weather quantity takes the form of a discrete parametric law. Recent studies demonstrated that the application of classification algorithms can successfully improve the skill of such discrete forecasts; however, a frequently emerging issue is that certain spatial and/or temporal dependencies could be lost between marginals. Based on visibility ensemble forecasts of the European Centre for Medium-Range Weather Forecasts for 30 locations in Central Europe, we investigate whether the inclusion of Copernicus Atmosphere Monitoring Service (CAMS) predictions of the same weather quantity as an additional covariate could enhance the skill of the post-processing methods and whether it contributes to the successful integration of spatial dependence between marginals. Our study confirms that post-processed forecasts are substantially superior to raw and climatological predictions, and the utilization of CAMS forecasts provides a further significant enhancement both in the univariate and multivariate setup. We also demonstrate that post-processing significantly improves the predictions of low visibility events, which opens the door for aeronautical applications.

KEYWORDS

Copernicus Atmosphere Monitoring Service (CAMS), ensemble calibration, ensemble copula coupling, multivariate post-processing, Schaake shuffle, visibility

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2024 The Author(s). *Meteorological Applications* published by John Wiley & Sons Ltd on behalf of Royal Meteorological Society.

1 | INTRODUCTION

Due to its significant impact on a wide range of aspects of aeronautical meteorology, visibility is a critical parameter for weather forecasters and pilots. Visibility can be reduced because of snow, rain, haze, and fog; the latter is reported “when visibility is reduced by water droplets to less than 1000 m” (WMO, 2022). In particular, the landing and takeoff of aircraft are impacted significantly by low visibility, and at international airports, when adverse visibility and ceiling conditions are observed, special protocols called low visibility procedures (LVPs) are applied (see Boudala et al., 2012; Dietz et al., 2019; Roquelaure & Bergot, 2009). The implementation of LVPs reduces the airport capacity, causing delays or even cancellations. Beyond the previously mentioned contexts, visibility’s importance extends notably into the realms of agriculture and maritime transportation. Additionally, poor visibility conditions stand out as one of the most common causes of road accidents. The World Meteorological Organization (WMO) defines visibility as “the greatest distance at which a black object of suitable dimensions (located on the ground) can be seen and recognized when observed against the horizon sky” (WMO, 1992). Due to the mentioned reasons, it is crucial to create the most accurate visibility forecasts possible, a task that, fortunately, an increasing number of meteorological services have been undertaking in recent years.

In the present day, weather forecasts are crafted through the use of numerical weather prediction (NWP) models. These models utilize initial values obtained from the previous states of the atmosphere, and the perturbation of these initial conditions results in the creation of a set of forecasts, commonly referred to as ensemble predictions. Despite notable advancements, NWP systems face challenges in achieving complete bias-free forecasts for various weather parameters, even within short lead times. Visibility emerges as a parameter with comparatively weak forecast skill, primarily because, unlike temperature, wind speed, or precipitation accumulation, most NWP models do not explicitly model visibility. Consequently, visibility forecasts necessitate derivation from predictions of related quantities such as relative humidity or precipitation (Chmielecki & Raftery, 2011). Specifically, at the European Centre of Medium-Range Weather Forecasts (ECMWF), the visibility parameter was integrated into the ECMWF Integrated Forecasting System (IFS) on May 12, 2015 (ECMWF, 2021). This parameter utilizes model projections of water vapor, cloud, rain, and snow, along with climatological aerosol fields, to estimate the visibility that would be recorded by weather observers.

To correct the biases and lack of calibration in the forecasts, one possible direction is the application of statistical post-processing. In recent years, numerous post-processing techniques have been developed to calibrate ensemble forecasts, for a recent overview we refer to Vanitsem et al. (2021). From this wide range of methods, the straightforward yet potent parametric ensemble model output statistic (Gneiting et al., 2005) approach maps ensemble forecasts onto a probability distribution, providing a pathway for assessing uncertainties associated with potential outcomes. The commonly used Bayesian model averaging (BMA; Raftery et al., 2005) also offers a complete predictive distribution for the examined weather variable; however, by now the machine learning-based distributional regression network (Rasp & Lerch, 2018) method has become a rather popular parametric technique (see Schultz & Lerch, 2022). Another widespread approach to post-processing is capturing the predictive distribution of a weather quantity with the help of its quantiles (see Bremnes, 2019; Friederichs & Hense, 2007). Naturally, such nonparametric calibration is also frequently combined with machine learning techniques. As examples, one can mention the quantile regression forests (Taillardat et al., 2016) or the Bernstein quantile network (Bremnes, 2020).

The utility of univariate post-processing techniques is supported by numerous studies; however, even today, many experts examine the disadvantage that the correlations between marginal distributions often get lost when applying such methods. This can be relevant to connections between locations, predictions initialized at the same time, or the loss of dependencies between weather quantities. Multivariate post-processing aims to restore potentially lost correlation structures, and there are already numerous parametric or nonparametric models available for achieving these goals (see Schefzik & Möller, 2018). One option is to consider multivariate predictive distributions (see Baran & Möller, 2015); however, the drawback of such models is that due to the estimation of a large number of parameters, they can easily encounter numerical problems. In contrast, the ensemble copula coupling (ECC; Schefzik et al., 2013) represents a two-step approach that involves initially generating samples from univariate predictive distributions. Subsequently, these samples are reorganized into the raw ensemble rank order structure, thus, this procedure can be interpreted as the application of an empirical copula. The Schaake shuffle (Clark et al., 2004) is another method that functions based on a similar principle to ECC; nevertheless, past research has suggested that ECC frequently demonstrates slightly better predictive performance (see Chen et al., 2024; Lakatos et al., 2023). For a comprehensive comparison of

such multivariate post-processing methods, we refer to Lerch et al. (2020).

Visibility observations at synoptic observation (SYNOP) stations are typically reported in discrete values according to WMO guidelines, namely “100 to 5,000 m in steps of 100 m, 6 to 30 km in steps of 1 km, and 35 to 70 km in steps of 5 km” (WMO, 2018, Section 9.1.2). A much finer, but still discrete reporting scheme is used in the Meteorological Aerodrome Report (METAR), where the maximal reported visibility is 10 km, 50 m steps are considered up to 800 m, 100 m steps from 800 m to 5 km, and 1 km steps from 5 to 10 km (WMO, 2022). As a result, the post-processing of visibility forecasts is reduced to a classification task, where the predicted probabilities of the different classes comprise the forecast distribution in the form of a discrete probability mass function (PMF). A recent study (Baran & Lakatos, 2023) demonstrated that the implementation of multilayer perceptron neural network (MLP; Goodfellow et al., 2016) and proportional odds logistic regression (POLR; McCullagh, 1980) classifiers can significantly enhance the forecast skill of raw ECMWF visibility predictions. Note, that these two methods displayed superior forecast skills also in the realm of post-processing ensemble predictions of total cloud cover (Baran et al., 2021). However, in contrast to Baran et al. (2021), the investigation of Baran and Lakatos (2023) did not include the exploration of the use of additional covariates that are not exclusively based on ECMWF forecasts of the target weather quantity. A natural first step in this direction is to consider further predictions of visibility that are independent of the operational ECMWF ensemble.

The Copernicus Atmosphere Monitoring Service (CAMS) is dedicated to monitoring and predicting the composition of Earth’s atmosphere. It provides information on various atmospheric components, including air quality, greenhouse gases, and aerosols (Pitkänen et al., 2023). CAMS delivers deterministic forecasts related to these components, generated by extending the IFS model developed by ECMWF with additional modules.

In this paper, we study the predictive accuracy of POLR and MLP techniques for calibrating ECMWF visibility ensemble forecasts, incorporating CAMS forecasts as covariates, in the contexts of both univariate and multivariate post-processing. Raw ECMWF visibility ensemble and climatology are considered as reference forecasts. The motivation is to provide spatially consistent calibrated forecasts of this complex weather quantity, which predictions might be well used, for instance, in route planning. In the case of univariate models, a particular focus is on aeronautical applications such as predicting low visibility events, which aspect was not investigated in

Baran and Lakatos (2023). The aim is to demonstrate that improved predictions, for example, of LVP states, can be obtained without tailoring the post-processing methods to the corresponding specific visibility events, which approach is more flexible than the one presented in Dietz et al. (2019).

The structure of the paper is outlined as follows. In Section 2, a concise description of the visibility datasets under study is presented. Section 3 provides a review of the considered univariate and multivariate post-processing methods, along with details on the approaches used for training data selection and the verification tools considered. The outcomes of our case study are detailed in Section 4, followed by a brief discussion and conclusions in Section 5.

2 | DATA

The visibility ensemble forecasts employed in this study consist of the operational ECMWF control (CTRL) and the corresponding 50 interchangeable ensemble members (ENS) produced by the ECMWF IFS. Regarding the latter 50 members of the ensemble, it is important to note that their interchangeability stems from the statistical indistinguishability arising from the perturbations of the initial conditions of the NWP. As mentioned in the Introduction, while visibility is primarily reported as a continuous variable in meters, the WMO suggests reporting observations in discrete categories, resulting in 84 different values in total, which can be described by the set

$$\mathcal{Y} = \{0, 100, 200, \dots, 4900, 5000, 6000, 7000, \dots, 29000, 30000, 35000, 40000, \dots, 65000, 70000\},$$

where the alignment of forecasts (provided in 1 m increments) with observations is achieved by rounding down to the nearest reported value.

As discussed, our goal is to show that incorporating CAMS visibility forecasts as an additional predictor in the post-processing methods described in Section 3—representing an independent yet comparable meteorological parameter—can significantly enhance the forecast skill of both the raw ECMWF predictions and the corresponding calibrated forecasts. The deterministic CAMS predictions are part of the global forecasts for atmospheric composition provided by the Copernicus Atmosphere Monitoring Service, a service provided by the ECMWF as part of the Copernicus program. The CAMS Global atmospheric composition forecast utilizes

the IFS, the same model employed for ECMWF weather forecasts. However, specific modules within CAMS have been enabled to account for aerosols, reactive gases, and greenhouse gases, which have been developed uniquely within the CAMS framework.

The ECMWF ensemble, the CAMS forecasts, and the corresponding observations span 2 years from January 2020 to December 2021, for 30 SYNOP stations in Germany, the Czech Republic, and Poland (refer to Figure 1). All forecasts are initialized at 0000 UTC, and we consider 20 different lead times ranging from 6 to 120 h, with a time step of 6 h. Notably, this dataset is fairly complete, with no missing observations.

The distribution of reported visibility observations is depicted in Figure 2. Approximately 80.78% of the total of 87,720 observations was not less than 10,000 m, while the proportions of low visibility events with observed visibility below 200 m (dense fog; Bari et al., 2016), 500 m, 1000 m (fog), and 3000 m are 3.19%, 5.17%, 6.24%, and 8.94%, respectively. Note that these thresholds are selected to match the ones in some recent studies on predicting low visibility events that have importance in aeronautical meteorology, for example, Penov and Guerova (2023), Shankar and Sahana (2023), He et al. (2024).

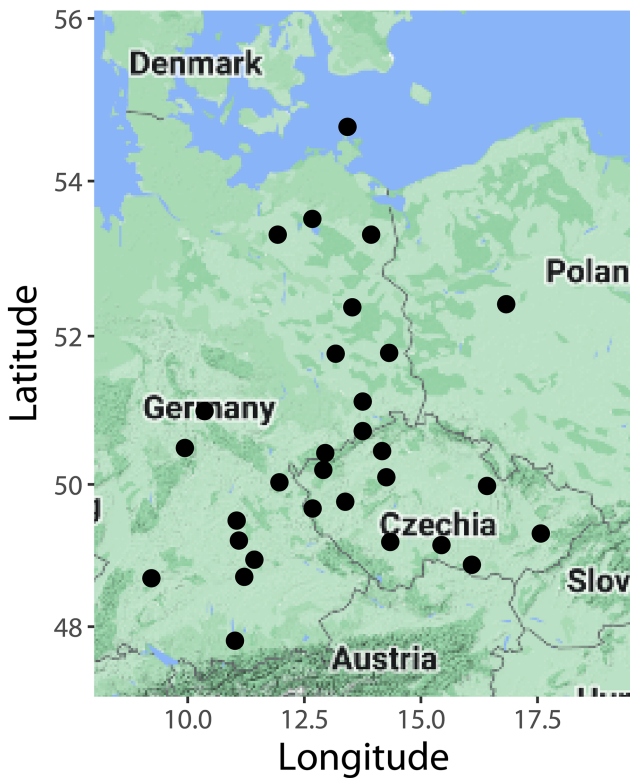


FIGURE 1 Locations of SYNOP observation stations corresponding to the ECMWF- and CAMS forecasts for calendar years 2020–2021.

3 | POST-PROCESSING- AND FORECAST EVALUATION METHODS

In what follows, let $\mathbf{f}^{(d)} = (f_1^{(d)}, f_2^{(d)}, \dots, f_{51}^{(d)})$ denote a 51-member ECMWF ensemble forecast for SYNOP station d ($d = 1, 2, \dots, 30$) for a given time point with a given lead time, where $f_1^{(d)} = f_{\text{CTRL}}^{(d)}$ represents the control run, and $f_2^{(d)}, f_3^{(d)}, \dots, f_{51}^{(d)}$ correspond to the 50 exchangeable members $f_{\text{ENS},1}^{(d)}, f_{\text{ENS},2}^{(d)}, \dots, f_{\text{ENS},50}^{(d)}$. Furthermore, let $f_0^{(d)} = f_{\text{CAMS}}^{(d)}$ be the matching CAMS prediction, while $Y^{(d)} \in \mathcal{Y} = \{y_1, y_2, \dots, y_{84}\}$ denotes the corresponding observed visibility with \mathcal{Y} defined in Section 2. Finally, station-specific data are combined into multivariate forecasts $\mathbf{f}_0, \mathbf{f}_1, \dots, \mathbf{f}_{51}$ and observations \mathbf{Y} , where $\mathbf{f}_\ell = (f_\ell^{(1)}, f_\ell^{(2)}, \dots, f_\ell^{(30)})^\top$, $\ell = 0, 1, \dots, 51$, and $\mathbf{Y} = (Y^{(1)}, Y^{(2)}, \dots, Y^{(30)})^\top$.

3.1 | Univariate post-processing

To simplify the presentation of calibration methods for marginals, which are identical to the approaches studied in Baran and Lakatos (2023), in this section we omit the indication of the SYNOP station and use notations $\mathbf{f} = \mathbf{f}^{(d)}$ and $Y = Y^{(d)}$ for forecasts and observations, respectively.

As Y is a discrete random variable, the predictive distribution of Y is specified by conditional probabilities

$$P(Y = y_k | \mathbf{f}), \quad k = 1, 2, \dots, 84, \quad (1)$$

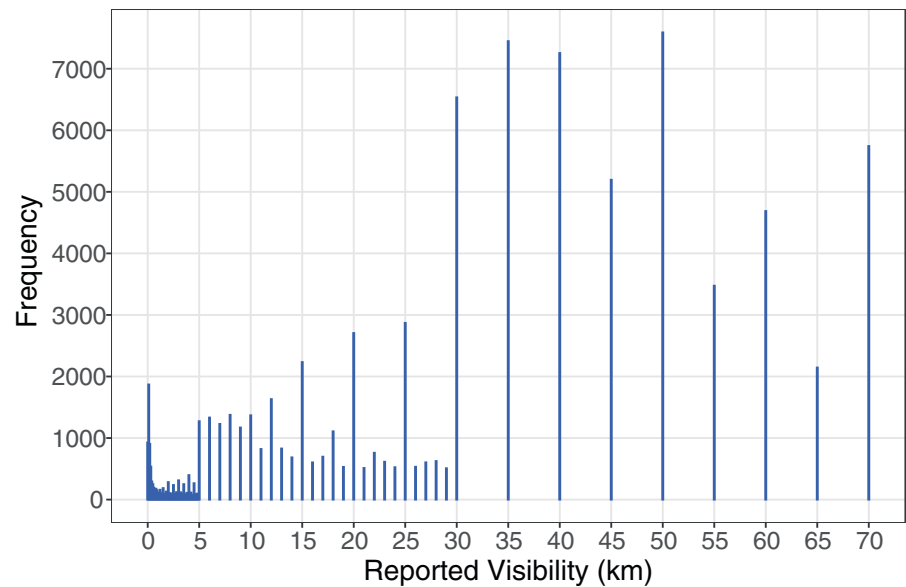
forming a conditional PMF with respect to the ensemble forecast.

Naturally, in (1) the raw forecast can be replaced by any feature vector \mathbf{x} derived from the ensemble and/or other covariates. These covariates can be location- or time-specific variables or additional forecasts of visibility or other related weather quantities.

3.1.1 | Multilayer perceptron neural network

In recent years, the application of various types of artificial neural networks (ANNs) for the calibration of ensemble forecasts has become increasingly popular. These networks have a significant advantage compared to classical parametric or nonparametric approaches, as they are more flexible, for example, in representing nonlinear relations between the ensemble forecasts and the parameters of the predictive distributions and in incorporating new explanatory variables in the modeling process.

FIGURE 2 Distribution of reported visibility observations for calendar years 2020–2021.



Multilayer perceptron (MLP) networks are a powerful class of ANNs widely employed for classification tasks in machine learning. A conventional MLP comprises multiple layers and nodes (neurons), where the value of each node is a result of the transformed (via an activation function) weighted sum of values from the nodes of the preceding layer, along with an added bias term. Input features are introduced to the input layer, and predictions regarding the distribution of different classes are generated in the output layer. The level of abstraction can be determined by tuning parameters, such as the number of hidden layers and the neurons within them. For more details, please refer to the work by Goodfellow et al. (2016).

3.1.2 | Proportional odds logistic regression

Proportional odds logistic regression (POLR) is an efficient parametric classification method for ordered classes, such as the visibility observations at hand. It specifies the conditional cumulative distribution of observed visibility Y given an M -dimensional feature vector \mathbf{x} as

$$P(Y \leq y_k | \mathbf{x}) = \frac{e^{\mathcal{L}_k(\mathbf{x})}}{1 + e^{\mathcal{L}_k(\mathbf{x})}}, \quad \text{where}$$

$$\mathcal{L}_k(\mathbf{x}) := a_k + \mathbf{x}^\top \mathbf{b}, \quad k = 1, 2, \dots, 84,$$

with $a_k \in \mathbb{R}$, $\mathbf{b} \in \mathbb{R}^M$, and $a_1 < a_2 < \dots < a_{84}$. Hence, POLR modeling of visibility classes results in a total of $84 + M$ unknown parameters to be estimated.

3.2 | Multivariate methods

The purpose of applying univariate post-processing methods is to fine-tune and calibrate the forecasts predicted by the NWP model, addressing various accuracy issues or biases. However, when applying these techniques, there is a risk of losing the spatial, temporal, or even inter-variable correlations between marginals. Here, we consider two easy-to-implement yet powerful two-step techniques that can correct this particular deficiency and restore dependence between marginal distributions (see Lakatos et al., 2023; Lerch et al., 2020). Both methods are based on reordering a sample generated from the calibrated predictive distribution obtained in the first step, where one can either consider equidistant quantiles or draw random or stratified samples. As Chen et al. (2024) reveal only minor differences in skill between the multivariate forecasts corresponding to the various sampling strategies, here we will work with the equidistant quantiles of the post-processed univariate forecasts.

3.2.1 | Ensemble copula coupling

Ensemble copula coupling (ECC) leverages the raw ensemble rank structure (with ties resolved at random) and capitalizes on the ample information embedded in ensemble predictions concerning dependencies. Following the notations of Lakatos et al. (2023) this iterative algorithm can be described as follows:

1. Generate a sample $\hat{f}_1^{(d)}, \hat{f}_2^{(d)}, \dots, \hat{f}_K^{(d)}$ for each dimension d , matching the size of the raw ensemble ($K = 51$),

from the calibrated marginal predictive distribution assumed to be sorted in ascending order.

2. Consider permutations

$$\boldsymbol{\pi}_d = (\pi_d(1), \pi_d(2), \dots, \pi_d(K))$$

of $\{1, 2, \dots, K\}$ corresponding to the rank order structure of the raw ensemble $f_1^{(d)}, f_2^{(d)}, \dots, f_K^{(d)}$, namely $\pi_d(k) := \text{rank}(f_k^{(d)})$, where ties are resolved at random. The ECC calibrated sample $\tilde{f}_1^{(d)}, \tilde{f}_2^{(d)}, \dots, \tilde{f}_K^{(d)}$ for location d is obtained by rearranging the sample generated in step 1 according to permutation $\boldsymbol{\pi}_d$, that is

$$\tilde{f}_k^{(d)} := f_{\pi_d(k)}^{(d)}, \quad k = 1, 2, \dots, K, \quad d = 1, 2, \dots, 30.$$

3.2.2 | Schaake shuffle

The other nonparametric multivariate post-processing technique we consider is the Schaake shuffle (SSh; Clark et al., 2004), which involves past observations to create a dependence template aimed at restoring correlations between the marginal distributions. Samples from the calibrated univariate predictive distributions are then rearranged to match the rank order structure of this selected set of historical observations of the corresponding cardinality. In this way, the SSh method offers the flexibility to generate post-processed forecasts of various sizes, provided there is a sufficiently long historical climatological dataset available. However, to ensure a fair comparison with ECC methods, we restrict the sample size to that of the raw ensemble and use the same sampling strategy.

Note that the general approach is to choose the dependence template for a given forecast date from a similar year period, for instance, by randomly selecting only from dates of the past years within 30 calendar days around the actual one. However, due to the limitations of our visibility dataset—2 years of observations versus the 24 years of Clark et al. (2004)—here we consider randomly selected observations from the past, which approach, according to Lakatos et al. (2023), appeared successful in the case of wind speed and precipitation accumulation.

3.3 | Training data selection

Each statistical post-processing method, including both the MLP and POLR models introduced in Section 3.1, conducts parameter estimation using training data, which comprises a collection of prior forecasts and

observations predating the current forecast time. Various approaches exist for selecting this training dataset in terms of both spatial and temporal decomposition.

A frequently used method of spatial selection, known as local estimation (Thorarinsdottir & Gneiting, 2010), entails using forecasts and observations specific to the station under consideration to estimate model parameters or weights. This approach could offer significant advantages as it effectively integrates the station's characteristics and local conditions into the calibration process. Another widely used and quite powerful alternative to the latter selection process is to treat together historical data of all stations of the ensemble domain, which, for a short training period, can be more efficient than local modeling. Nevertheless, this method, referred to as regional, is not particularly suitable for expansive and diverse domains, since it may struggle to effectively capture the varied complexities and nuances present across such landscapes. Due to the factors mentioned, in the case of regional modeling, all locations share the same set of parameters on a specific day to derive the parameters of the predictive distribution. A middle ground between regional and local methods is semi-local modeling, where stations with similar characteristics are clustered together, and regional modeling is applied within these clusters (Lerch & Baran, 2017). Similar to Baran and Lakatos (2023) and Baran and Baran (2024), in the present study all three approaches to spatial selection of training data are considered.

The efficiency of post-processing methods is further influenced by the length of the training period and where the training period is positioned in time. One commonly used technique for temporal selection is when the training data is selected using a so-called sliding window approach, where the data for the preceding n calendar days before the actual forecast date are chosen. Alternatively, fixing the training period to a specific extensive time frame in the past can also be suitable for implementing the modeling process, especially in the case of machine learning-based techniques requiring large training datasets. Both approaches ensure that post-processed forecasts are verified against unseen data. For a recent comparison of various time-adaptive training schemes, we refer to Lang et al. (2020).

3.4 | Verification scores

Following the suggestions of Gneiting and Raftery (2007), the predictive performance of the studied univariate and multivariate visibility forecasts is evaluated with the help of proper scoring rules, as they simultaneously address both the calibration and sharpness of the predictive

distribution. To be consistent with Baran and Lakatos (2023), in the univariate case, we consider the logarithmic score (LogS; Good, 1952) and the continuous ranked probability score (CRPS; Wilks, 2019, Section 9.5.1). For a discrete visibility forecast F represented by a PMF $p_F(y)$ on \mathcal{Y} the former is defined as

$$\text{LogS}(F, y) := -\log(p_F(y)),$$

whereas the latter equals

$$\begin{aligned} \text{CRPS}(F, y) &= \sum_{k=1}^{84} p_F(y_k) |y_k - y| \\ &- \sum_{k=2}^{84} \sum_{\ell=1}^{k-1} p_F(y_k) p_F(y_\ell) |y_k - y_\ell|, \end{aligned} \quad (2)$$

which is the discrete version of the general representation

$$\text{CRPS}(F, y) = \mathbb{E} |Y - y| - \frac{1}{2} \mathbb{E} |Y - Y'| \quad (3)$$

derived by Gneiting and Raftery (2007), where Y and Y' are independent random variables with finite second moments distributed according to F . Note that the LogS is local, as it uses just the predicted probability of the materialized observation, whereas the CRPS considers skillful forecasts where values with high probabilities are close to the observation.

Furthermore, we also consider the Brier score (BS; Wilks, 2019, Section 9.4.2) to quantify the predictive performance of the competing discrete forecasts with respect to events that the observed visibility y is in a set (interval) \mathfrak{F} . Given again a PMF $p_F(y)$ on \mathcal{Y} representing the visibility forecast F , the BS is defined as

$$\text{BS}(F, y; \mathfrak{F}) := \left[\sum_{k: y_k \in \mathfrak{F}} p_F(y_k) - \mathbb{I}_{\{y \in \mathfrak{F}\}} \right]^2,$$

where \mathbb{I}_H denotes the indicator function of an event H .

Form (3) allows a direct extension of the CRPS to multivariate forecasts. For a D -dimensional predictive cumulative distribution function (CDF) F and vector $\mathbf{y} = (y^{(1)}, y^{(2)}, \dots, y^{(D)})^\top$, the energy score (ES; Gneiting & Raftery, 2007) is defined as

$$\text{ES}(F, \mathbf{y}) = \mathbb{E} \|\mathbf{Y} - \mathbf{y}\| - \frac{1}{2} \mathbb{E} \|\mathbf{Y} - \mathbf{Y}'\|,$$

where now \mathbf{Y} and \mathbf{Y}' are independent random vectors having distribution F , and $\|\cdot\|$ denotes the Euclidean

distance. For a K -member forecast ensemble or sample drawn from the predictive distribution, one has to consider the empirical CDF F_K resulting in the ensemble energy score (Gneiting et al., 2008)

$$\text{ES}(F_K, \mathbf{y}) = \frac{1}{K} \sum_{j=1}^K \|\mathbf{f}_j - \mathbf{y}\| - \frac{1}{2K^2} \sum_{j=1}^K \sum_{k=1}^K \|\mathbf{f}_j - \mathbf{f}_k\|.$$

Finally, we also report the (ensemble) variogram score of order p (VS^p ; Scheuerer & Hamill, 2015), which is more sensitive to the errors in the specification of correlations than the ES. Given an ensemble forecast

$\mathbf{f}_k = (f_k^{(1)}, f_k^{(2)}, \dots, f_k^{(D)})^\top$, $k = 1, 2, \dots, K$, it equals

$$\text{VS}^p(F_K, \mathbf{y}) = \sum_{i=1}^D \sum_{j=1}^D \omega_{ij} \left(|y^{(i)} - y^{(j)}|^p - \frac{1}{K} \sum_{k=1}^K |f_k^{(i)} - f_k^{(j)}|^p \right)^2 \quad (4)$$

where $\omega_{ij} \geq 0$ is the weight for coordinate pair (i, j) . Here, ω_{ij} is the negative exponential of the geographical distance between the corresponding SYNOP stations i and j expressed in 100 km. The usual choices for order p are 0.5 and 1, in our case study we consider the former and use the notation VS for $\text{VS}^{0.5}$. Note that all five scoring rules defined above are negatively oriented—that is, smaller values are better—and all, except the BS, are implemented in the R package `scoringRules` (Jordan et al., 2019).

In the case study of Section 4 the predictive performance of a forecast F with a given forecast horizon in terms of a score \mathcal{S}_F is quantified by the mean score value $\bar{\mathcal{S}}_F$ over all forecast cases used for verification. Furthermore, we also consider skill scores (Murphy, 1973)

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

providing the relative improvement of a forecast F in terms of the score \mathcal{S}_F with respect to a reference forecast F_{ref} resulting in a mean score value $\bar{\mathcal{S}}_{F_{\text{ref}}}$. Skill scores are positively oriented (the larger the better) and often reported in percentage values. In Section 4 we investigate continuous ranked probability skill score (CRPSS), logarithmic skill score (LogSS), Brier skill score (BSS), energy skill score (ESS), and variogram skill score (VSS).

In addition to reporting the various score and skill score values of the competing visibility forecasts, we also address the significance of the score differences and uncertainty in score values by providing 95% confidence bounds. These confidence intervals are calculated using

2000 block-bootstrap samples according to the stationary bootstrap scheme with the mean block length formula specified by Politis and Romano (1994).

Calibration of univariate ensemble forecasts can also be assessed graphically with the help of verification rank histograms (Wilks, 2019, Section 9.7.1). For a properly calibrated K -member ensemble, the ranks of the verifying observations with respect to the corresponding forecasts follow a discrete uniform distribution on the set $\{1, 2, \dots, K+1\}$, and the shape of the corresponding histogram reflects the source of the lack of calibration. Multivariate generalizations of the verification rank are based on the concept of pre-ranks and the corresponding histograms display the pre-rank of the observation with respect to the ensemble forecast with ties resolved at random (Gneiting et al., 2008). Here, we consider the average and band-depth ranks suggested by Thorarinsdottir et al. (2016) and energy score and dependence (vario-gram) ranks studied in detail, for example, by Allen et al. (2024). The average rank is simply the mean of the univariate ranks and the interpretation of the corresponding histogram is similar to the univariate case: \cup - and \cap -shaped histograms refer to under- and overdispersion, respectively, whereas biased forecasts result in triangular shapes. Band-depth histograms are based on the discrete version of ordering multivariate functional data proposed by López-Pintado and Romo (2009), while energy score and dependence histograms use the idea of Knüppel et al. (2022) of proper score-based pre-ranks. Right-skewed band-depth histograms indicate overdispersion, left-skewed ones underdispersion or bias, whereas \cup - and \cap -shaped histograms mean too low- and high correlations in the ensemble. Furthermore, low score-based ranks indicate that the observation is similar to the ensemble members in terms of the given scoring rule, whereas outliers result in high pre-ranks. Note, that in the case of the dependence ranking we consider the same geographical distance dependent weights for the vario-gram score as in (4). Finally, as a measure of the deviation of verification ranks from uniformity we consider the reliability index (RI; Delle Monache et al., 2006) defined as

$$\text{RI} := \sum_{r=1}^{K+1} \left| \rho_r - \frac{1}{K+1} \right|, \quad (5)$$

where ρ_r is the relative frequency of rank r .

4 | RESULTS

We first explore the advantage of applying CAMS visibility forecasts as additional covariates in POLR and MLP

methods for discrete post-processing of ECMWF ensemble predictions of visibility compared to models based on the raw ECMWF ensemble only. We also study the predictive performance of the corresponding multivariate forecasts obtained with the help of ensemble copula coupling and Schaake shuffle. To be consistent with the results of Baran and Lakatos (2023), all investigated MLP and POLR classifiers rely on 350-day rolling training periods, and each lead time is treated separately. We consider local (MLP-L and POLR-L), regional (MLP-R and POLR-R) and clustering-based semi-local (MLP-C and POLR-C) training with four clusters derived using k -means clustering of SYNOP stations with respect to three-dimensional feature vectors consisting of observed frequencies of visibility intervals 0–5000, 5000–30,000, and 30,000–70,000 m in the training period. Both univariate and multivariate predictions are validated on forecast-observation pairs for calendar year 2021 and as reference forecasts, we consider the raw ECMWF visibility ensemble and 30-day climatology.

The reference MLP and POLR models studied by Baran and Lakatos (2023) are based on 8D input feature vectors

$$\left(\tilde{f}_{\text{CTRL}}, \bar{f}_{\text{ENS}}, s^2, p_1, p_2, p_3, \beta_1, \beta_2 \right)^\top, \quad (6)$$

where $\tilde{f}_{\text{CTRL}} := f_{\text{CTRL}}/70000$ is the normalized control member, \bar{f}_{ENS} is the mean of the 50 (normalized) exchangeable members, s^2 is the variance of the 51-member (normalized) operational ensemble, p_1 , p_2 and p_3 are the proportions of ensemble members predicting visibility up to 1000 m, 1000–2000 m and more than 30,000 m, respectively, whereas β_1 and β_2 are annual base functions addressing seasonal variations defined as

$$\beta_1(d) := \sin(2\pi d/365) \quad \text{and} \quad \beta_2(d) := \cos(2\pi d/365),$$

where d denotes the day of the year (see Dabernig et al., 2017). We note that normalization is applied to stabilize modeling numerically, as it makes all input features commensurate.

Here, we add the normalized CAMS forecast $\tilde{f}_{\text{CAM5}} := f_{\text{CAM5}}/70000$ to the input features and investigate the forecast skill of CAMS extended local, regional, and semi-local MLP and POLR post-processing based on the feature vector

$$\left(\tilde{f}_{\text{CTRL}}, \bar{f}_{\text{ENS}}, \tilde{f}_{\text{CAM5}}, s^2, p_1, p_2, p_3, \beta_1, \beta_2 \right)^\top \quad (7)$$

together with the corresponding multivariate forecasts derived with the help of the two-step approaches presented in Section 3.2.

4.1 | Univariate performance

The predictive performance of univariate visibility forecasts is evaluated with the help of the mean CRPS and mean LogS over all forecast cases in the verification period, together with the corresponding skill scores. As in Baran and Lakatos (2023), extremely low predicted probabilities resulting in numerical issues in LogS calculations are handled by introducing a probability threshold $p_{\min} = 2.75 \times 10^{-5}$, which replaces all values in the predictive PMFs which are below this probability. This particular choice of p_{\min} ensures that with a 1% probability the corresponding reported visibility materializes at least once in a year; however, it is low enough, so there are no observable differences in terms of other scores between the original and the adjusted and renormalized PMFs. Furthermore, as for the reference POLR approaches based on feature vector (6) the weights of \tilde{f}_{CTRL} and \tilde{f}_{ENS} are forced to be non-negative, we apply the same constraint also for the coefficient of \tilde{f}_{CAMS} of the CAMS extended local, semi-local and regional POLR models. For further details of the implementation of MLP and POLR classifiers see Baran and Lakatos (2023) and Hemri et al. (2016).

As demonstrated in Section 4.2 of Baran and Lakatos (2023), in terms of the mean CRPS and mean LogS, up to 120 h, all MLP and POLR post-processed forecasts have shown significant improvement compared to both the raw ensemble and climatology. The extent of this improvement is particularly striking in the case of the raw ensemble, which we have chosen to exclude from our univariate analysis. Figure 3a displays the mean CRPS of climatological and CAMS extended MLP and POLR forecasts as functions of the forecast horizon. All post-processed predictions outperform climatology for

all lead times, and the ranking of the various approaches is identical to the one for the corresponding reference models based on feature vector (6), see Figure 2 of Baran and Lakatos (2023). The advantage achieved in mean CRPS by utilizing the CAMS extended feature vector (7) over models based solely on ECMWF forecasts is illustrated by the CRPS values of Figure 3b. In general, it can be observed that the integration of CAMS into the post-processing workflow for regional models yields the lowest benefits, but this advantage is still noteworthy. Note that the best performing POLR and MLP variants, which are the local POLR model and the semi-local MLP, profit the most from the extension of the input features with an average CRPS of around 3%. One should also remark that the improvement in mean CRPS from including CAMS predictions might further decrease for forecast horizons beyond 120 h.

The mean LogS values of Figure 4a tell the same story about the various post-processing approaches as Figure 3 of Baran and Lakatos (2023). Adding the CAMS forecast to the input features of the local, semi-local, and regional MLP and POLR models does not change the ranking of the forecasts; however, as portrayed in Figure 4b, it results in a modest 0.4–0.9% average improvement in LogS.

To address the significance of gain of CAMS extended post-processing models with respect to the corresponding forecasts based only on ECMWF predictions, in Figure 5 we accompany the CRPS and LogSS of the MLP-C and POLR-L approaches, which are the best model variants according to Baran and Lakatos (2023), by 95% block-bootstrap confidence intervals. As depicted in Figure 5a,b, up to the maximal studied lead time of 120 h, the improvement in terms of both scores is significant at a 5% level.

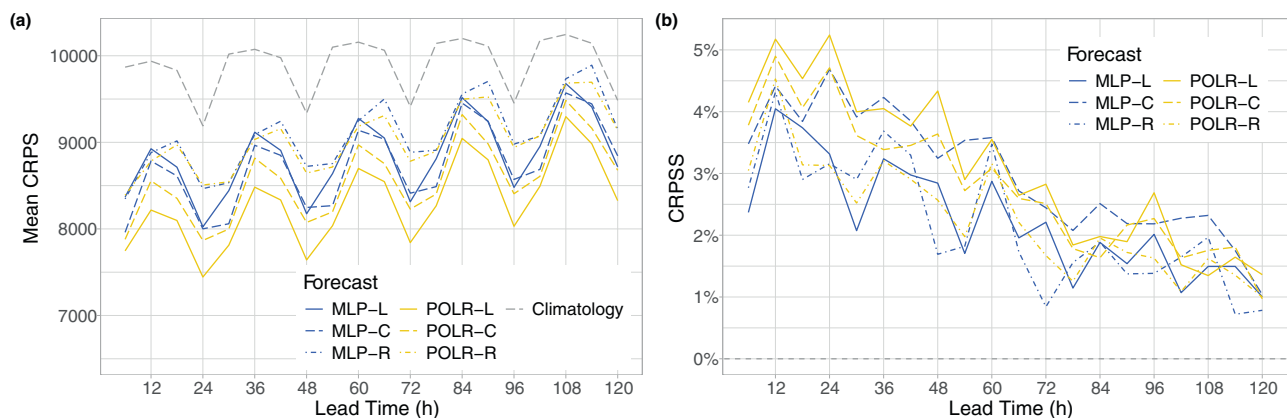


FIGURE 3 Mean CRPS of climatological and post-processed visibility forecasts based on CAMS extended MLP and POLR models (a) and CRPS of CAMS extended post-processed forecasts with respect to the corresponding MLP and POLR models based only on ECMWF predictions (b) as functions of the lead time.

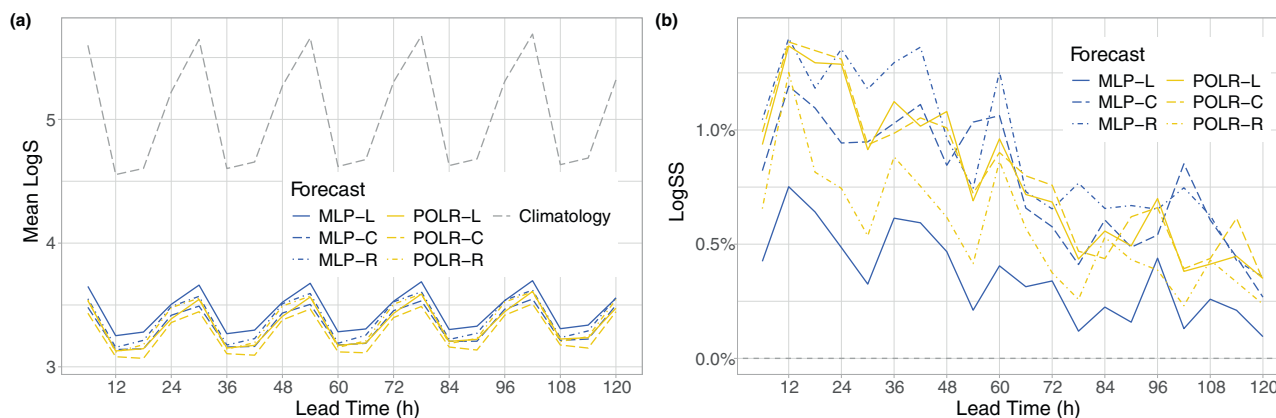


FIGURE 4 Mean LogS of climatological and post-processed visibility forecasts based on CAMS extended MLP and POLR models (a) and LogSS of CAMS extended post-processed forecasts with respect to the corresponding MLP and POLR models based only on ECMWF predictions (b) as functions of the lead time.

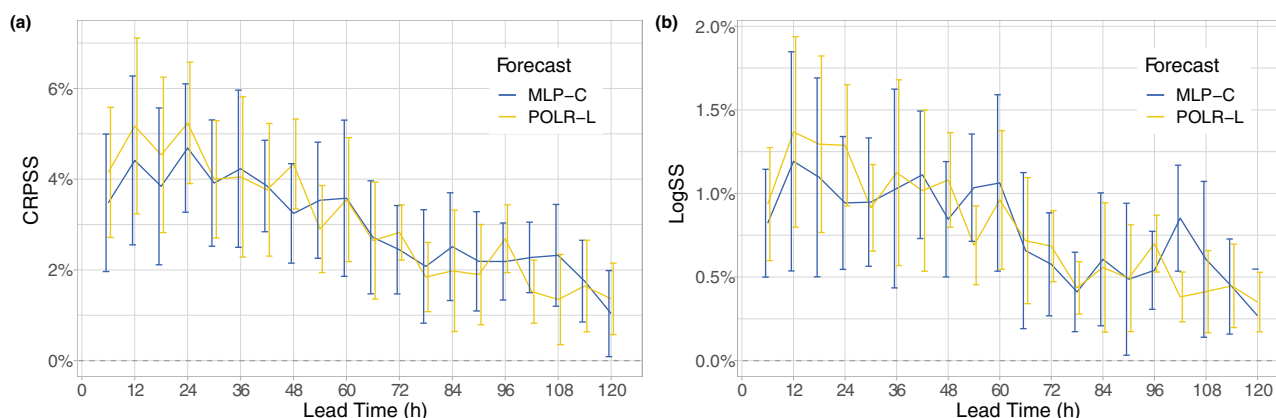


FIGURE 5 CRPSS (a) and LogSS (b) of CAMS extended MLP-C and POLR-L forecasts with respect to the corresponding models based only on ECMWF predictions together with 95% confidence bounds as functions of the lead time.

Finally, we investigate the skill of the best-performing CAMS extended MLP-C and POLR-L methods for the visibility range considered in aeronautical meteorology. We censor the observed classes from above at 10 km, so reported observations exceeding this value are rounded down to 10 km, resulting in 56 different values, while the corresponding probabilities are pooled. Naturally, in this case definition (2) of the CRPS should be modified accordingly.

As depicted in Figure 6a, censoring at 10 km does not change the ranking of the different forecasts in terms of the mean CRPS. However, compared to Figure 2 of Baran and Lakatos (2023), the advantage of climatology over the (CAMS extended) raw ensemble is much smaller, and the difference between climatological and post-processed forecasts is also reduced compared to Figure 3a. Nevertheless, the local POLR model significantly outperforms climatology for all lead times with a CRPSS averaging 17.7% (Figure 6b), whereas for the semi-local MLP

approach, the average gain in CRPSS over climatology is 9.5%, and the difference fails to be significant at a 5% level for forecasts corresponding to 1200 UTC.

The Brier skill scores of Figure 7 assessing the relative performance of post-processed forecasts with respect to climatology for various low visibility events discussed in Section 2 convey the same message. The advantage of the POLR-L approach is significant in all four cases for all lead times with BSS values between 9.3% and 21.4%, while the MLP-C method again cannot show a significantly positive skill for 1200 UTC observations.

4.2 | Multivariate post-processing

Continuing our analysis, we delve into the comparison of forecast skill among multivariate approaches detailed in Section 3.2. In the case of the applied multivariate methods, calibrated samples of size 51 are employed,

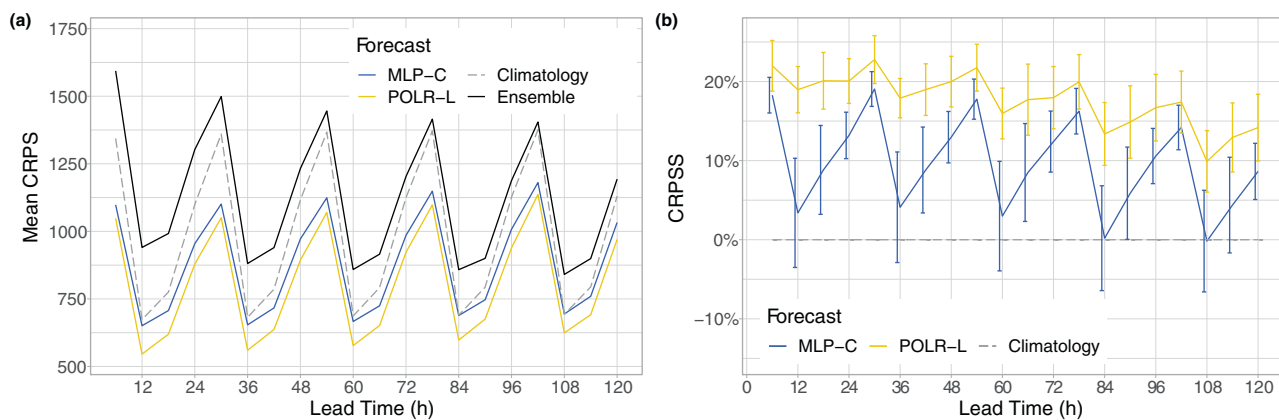


FIGURE 6 Mean CRPS of raw, climatological and CAMS extended MLP-C and POLR-L forecasts censored from above at 10 km (a) and CRPSS of CAMS extended post-processed forecasts with respect to climatology (b) together with 95% confidence bounds as functions of the lead time.

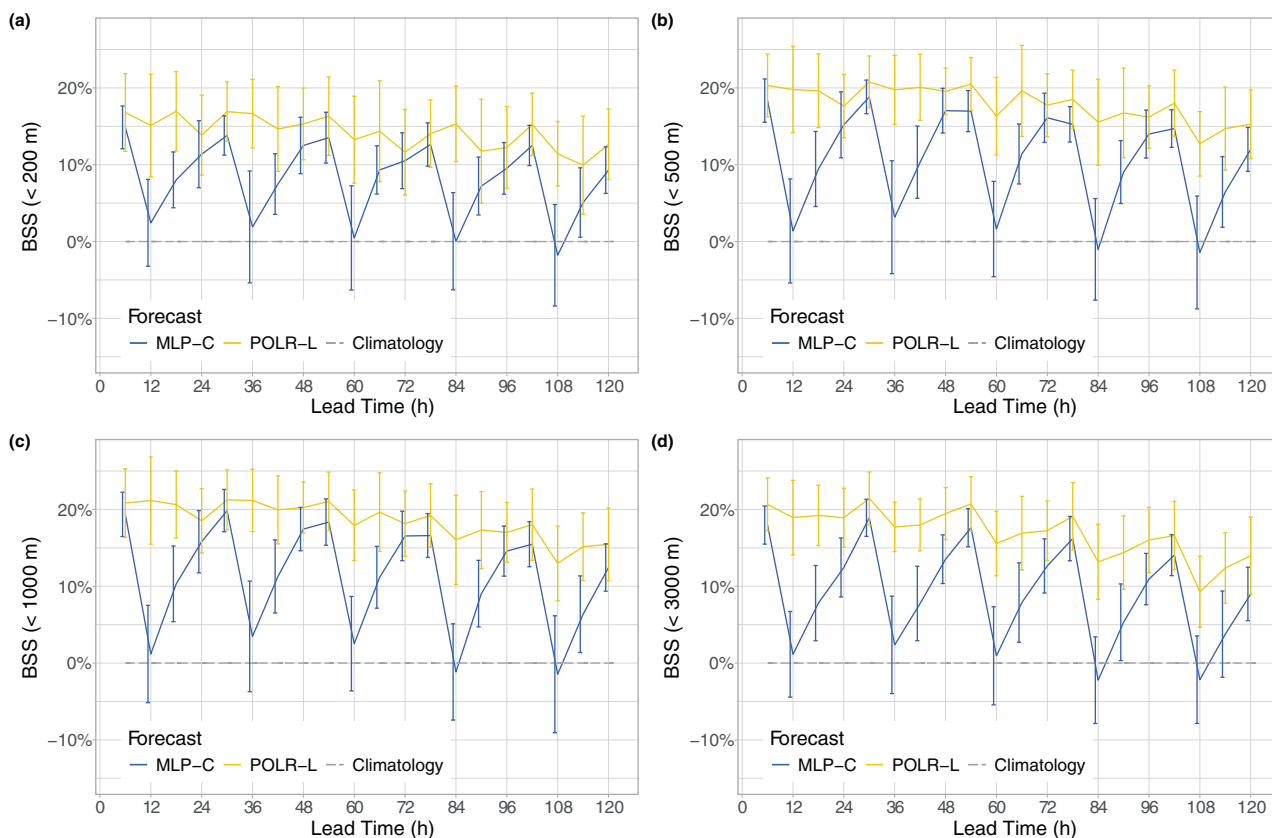


FIGURE 7 BSS of CAMS extended MLP-C and POLR-L forecasts with respect to climatology for the events where visibility is below 200 m (a), 500 m (b), 1000 m (c), and 3000 m (d) together with 95% confidence bounds as functions of the lead time.

mirroring the dimensionality of the raw ECMWF ensemble. When employing the SSh method, which relies on historical data to establish dependency templates, we utilize observations obtained from the rolling training window applied during univariate calibration. This decision is made based on the finding of Lakatos et al. (2023), that extending these pool dates to encompass the entire

available historical dataset may only yield negligible advantages. In our case, this entails selecting 51 days from a pool of 350 dates (which is the length of the training period). As mentioned, in both the ECC and SSh methods, the initial step entails generating samples from the calibrated PMFs, which we perform by considering their equidistant quantiles.

To simplify presentation, as univariate calibrated forecasts we consider the CAMS extended MLP-C and POLR-L predictions and their counterparts based only on the ECMWF ensemble. The corresponding multivariate predictions obtained with the help of the ECC and SSh are referred to as MLP-C ECC, POLR-L ECC, and MLP-C SSh, POLR-L SSh, respectively. We also report the performance of naive multivariate forecasts obtained by simply arranging the 51-member station-specific MLP-C and POLR-L samples into 30-dimensional vectors. In the multivariate context, notations MLP-C and POLR-L will refer to these “independent” predictions. Finally, as a multivariate climatological forecast, we consider the vector of 51 equidistant quantiles of 30-day climatological PMFs corresponding to the investigated SYNOP stations. This approach results in 51 continuous 30-dimensional vectors. In principle, one could also consider the vectors of station-specific 51-day climatological forecasts; however,

due to their discrete nature (all vector coordinates are elements of the discrete set \mathcal{Y}), they would substantially differ from the competing calibrated predictions.

First, we assess the performance of the previously discussed multivariate forecasts based on the energy score. Figure 8a clearly illustrates the substantial superiority of each CAMS extended multivariate forecast over the raw ensemble and the 30-day climatology in terms of this scoring rule. Notably, the semi-local MLP ECC and local POLR ECC exhibit the best predictive performance, extremely closely approaching their SSh counterparts, thereby minimizing performance disparities and highlighting the success of multivariate post-processing. According to the energy skill scores of Figure 8b, the univariate MLP-C and POLR-L models alone demonstrate a substantial average improvement of 12% over multivariate climatology, while the multivariate models show even more pronounced performance enhancements, averaging

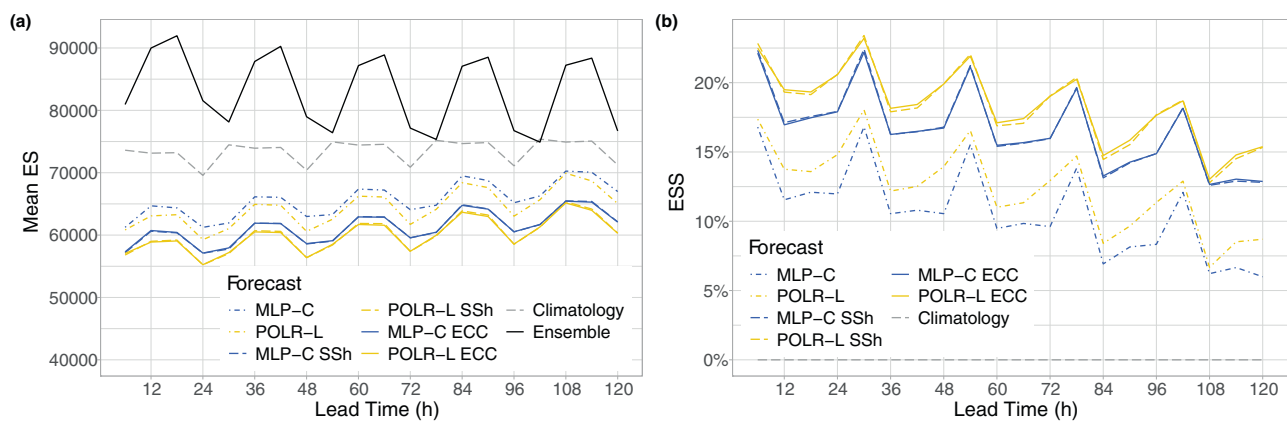


FIGURE 8 Mean ES of raw, climatological and post-processed multivariate visibility forecasts based on CAMS extended MLP-C and POLR-L models (a) and ESS of CAMS extended post-processed forecasts with respect to multivariate climatology (b) as functions of the lead time.

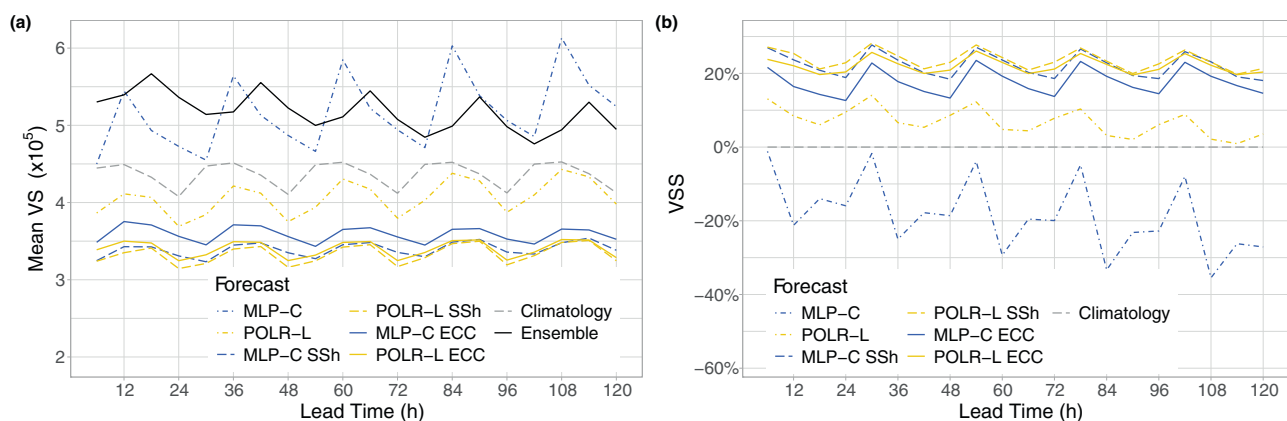


FIGURE 9 Mean VS of raw, climatological and post-processed multivariate visibility forecasts based on CAMS extended MLP-C and POLR-L models (a) and VSS of CAMS extended post-processed forecasts with respect to climatology (b) as functions of the lead time.

17%. This analysis emphasizes that, at least in our dataset, irrespective of the selected dependency template, multivariate models consistently exhibit improvement compared to independent calibration and climatology.

Figure 9a presents a slightly different picture in terms of model ranking, which might be a natural phenomenon

when considering the variogram score depicted here. This verification measure clearly illustrates the success of multivariate post-processing in the redefinition of covariance structures. As expected, (CAMS extended) naive models, such as MLP-C and POLR-L, do not capture the spatial dependencies and might underperform the raw

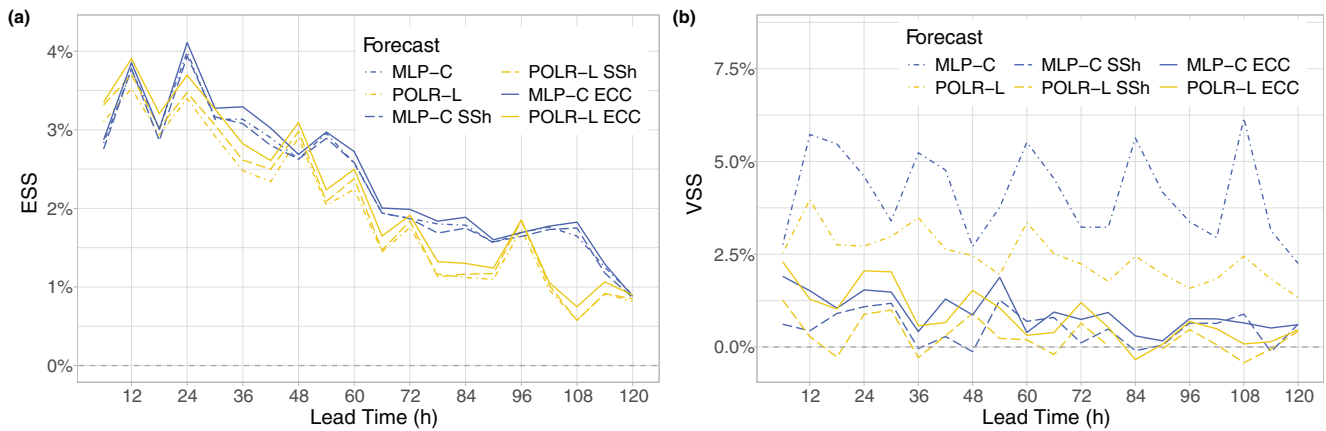


FIGURE 10 ESS (a) and VSS (b) post-processed multivariate visibility forecasts based on CAMS extended MLP-C and POLR-L models with respect to the corresponding forecasts based only on ECMWF predictions as functions of the lead time.

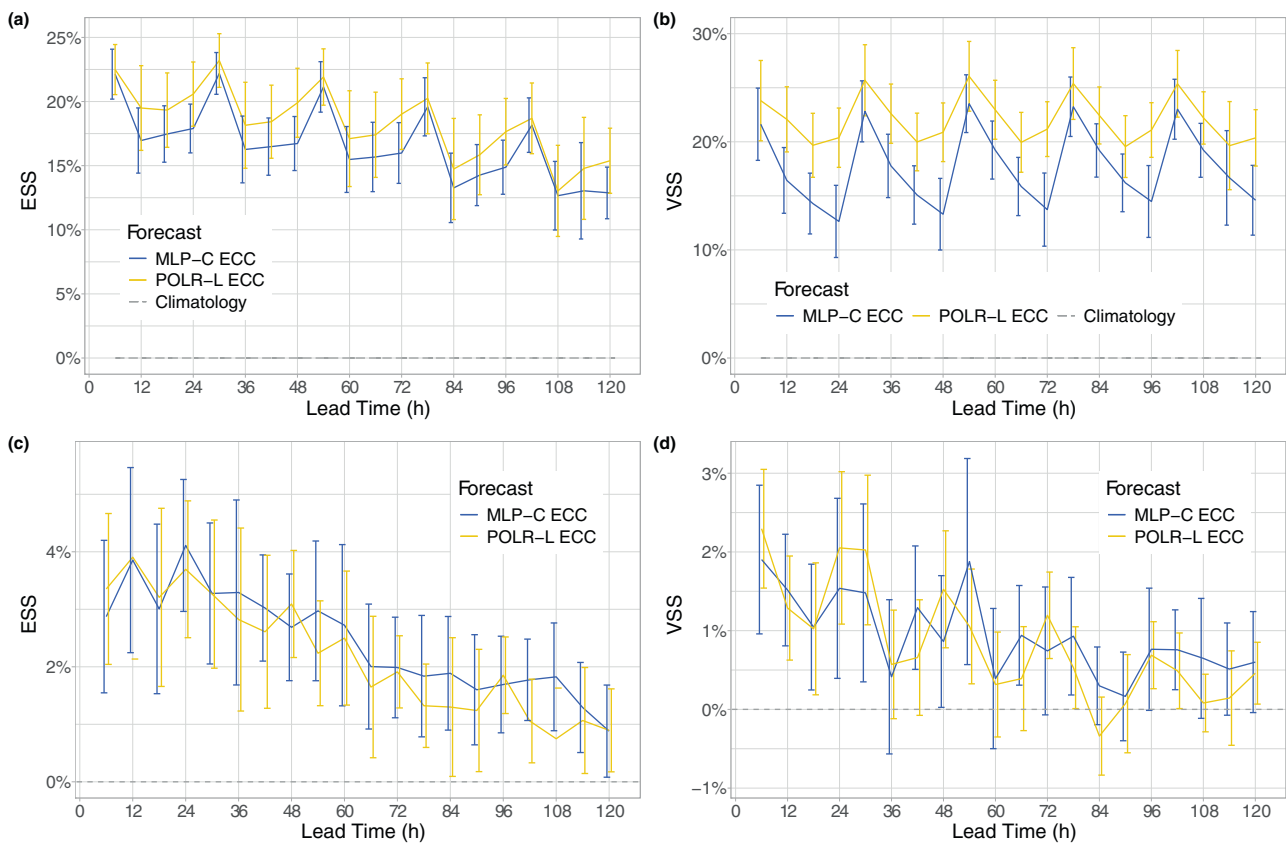


FIGURE 11 ESS (a,c) and VSS (b,d) of CAMS extended MLP-C ECC and POLR-L ECC multivariate forecasts with respect to multivariate climatology (a,b) and to the corresponding MLP-C ECC and POLR-L ECC forecasts based only on ECMWF predictions (c,d) together with 95% confidence bounds as functions of the lead time.

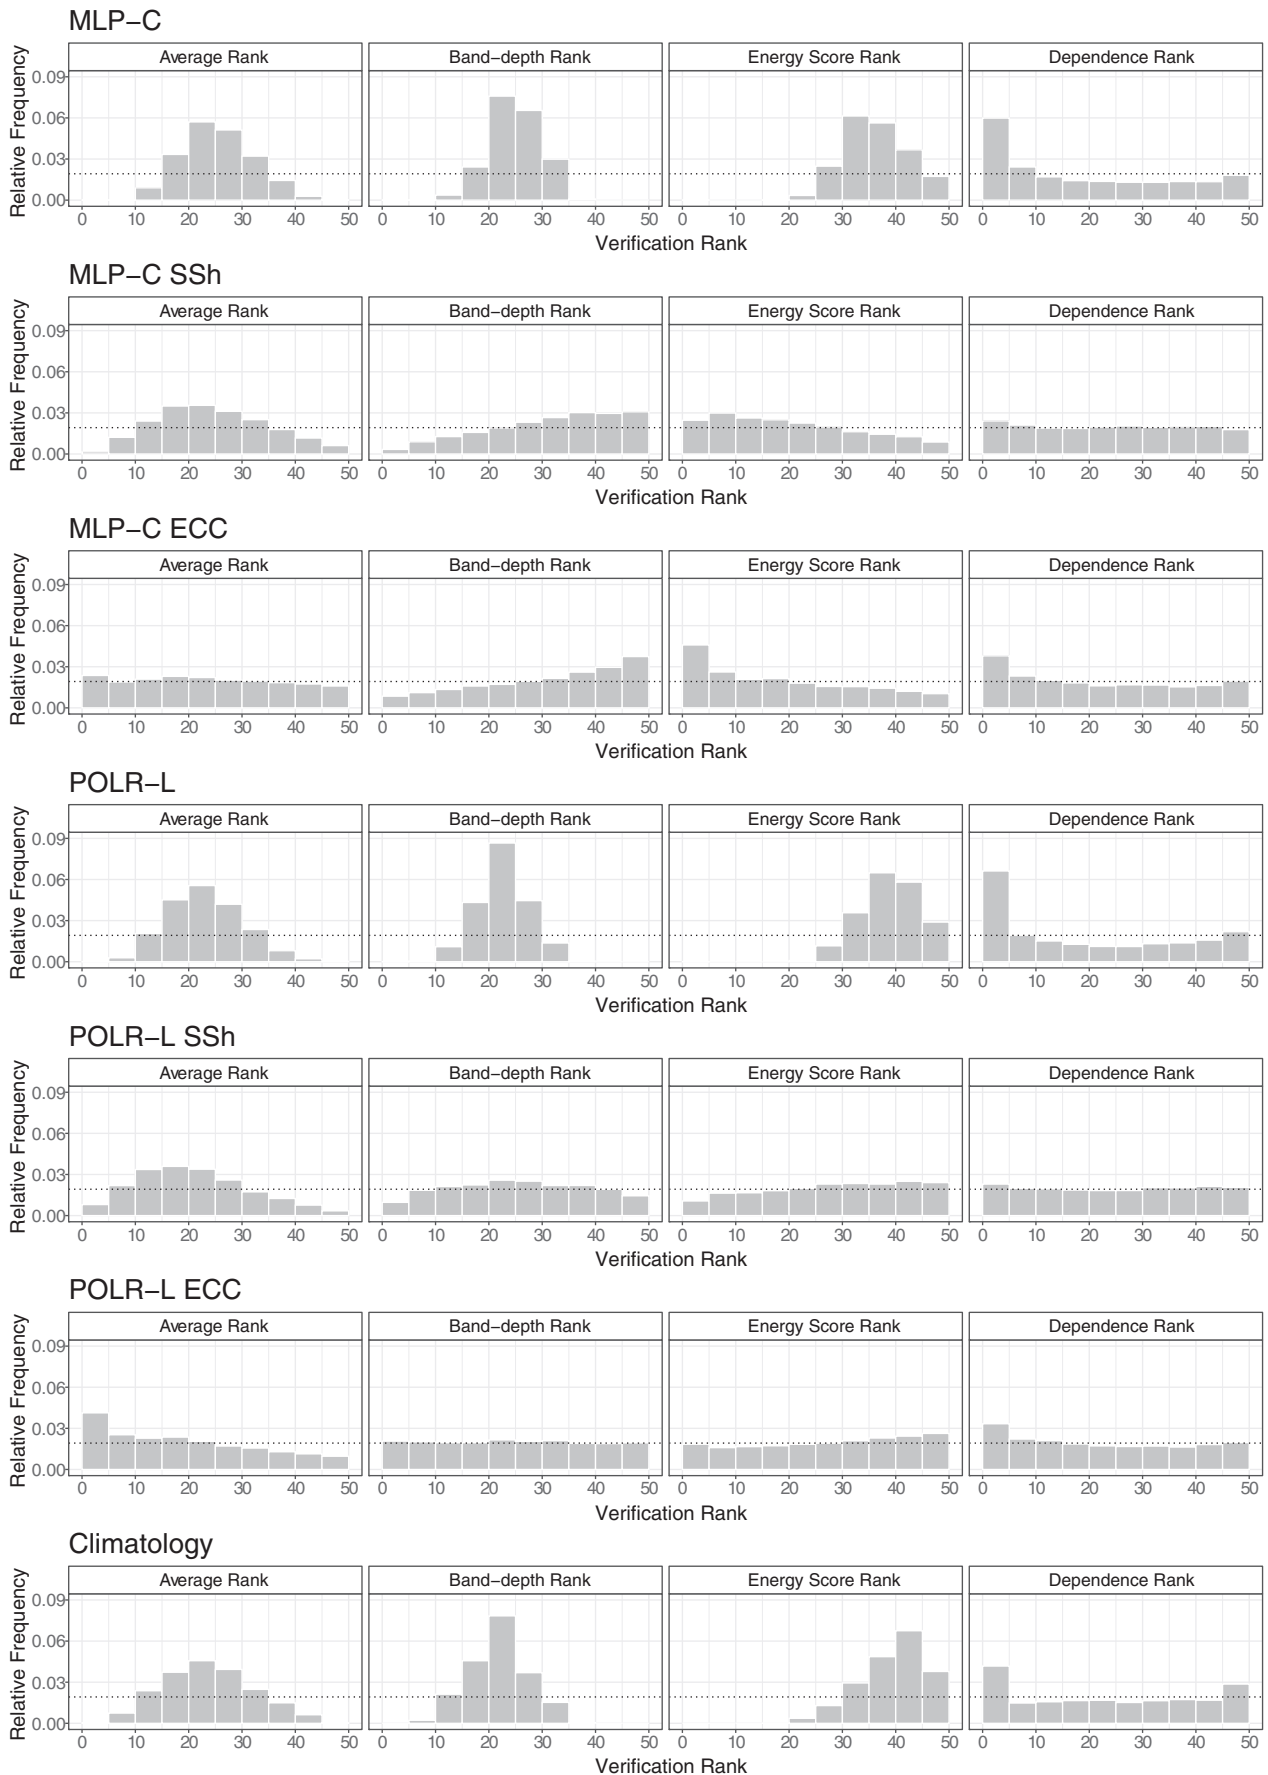


FIGURE 12 Rank histograms of climatological and post-processed multivariate visibility forecasts based on CAMS extended MLP-C and POLR-L models.

forecast vectors and multivariate climatology, as well. According to the VSS values of Figure 9b, the two-step multivariate models demonstrate a notable average improvement of over 17.6% with respect to the reference climatological predictions. Among these approaches, the POLR variants exhibit the highest skill, particularly the POLR-L SSh configuration. Moreover, it is noteworthy that a more pronounced advantage was attained with the SSh models obtaining dependence templates from past observations, compared to the ECC forecasts utilizing the rank structure of raw ensemble.

In Figure 10, we illustrate how the CAMS extended post-processed forecasts compare to the corresponding multivariate predictions from classifiers that ignore CAMS forecasts. Both the ESS and VSS reveal a diminishing advantage of CAMS integration with increasing forecast horizons. In the case of the ES the differences between the various post-processed forecasts are minor (see Figure 10a) with the MLP-C ECC variant showing the most overall improvement. The VSS values of Figure 10b again tell a different story. With independent MLP-C and POLR-L models, there is an average improvement of around 4% and 2.5%, respectively, in variogram scores suggesting that using CAMS as a feature helps in

better preserving spatial dependencies. For the multivariate models, the gain is much smaller, and there are several lead times where some models display even negative skill.

For better clarity, in Figure 11, we accompanied the ESS and VSS values of MLP-C ECC and POLR-L ECC forecasts – which are overall considered the best performers – with 95% bootstrap confidence bounds. In cases where climatology serves as the reference, as seen in Figure 11a,b, in general, the local POLR-based ECC outperforms its MLP-C counterpart. However, for forecasts corresponding to 0600 UTC, both in terms of ES and VS, the difference is not significant, indicating that the advantage of the POLR-L ECC approach over the MLP-C ECC is more apparent for the other three studied observation times (1200, 1800, 2400 UTC). Furthermore, Figure 11c,d demonstrate that incorporating CAMS forecasts as an additional feature leads to a significant enhancement in ES for all investigated lead times, whereas for VS there are a few instances of longer lead times exhibiting negative skill within the confidence bounds.

In addition to the previous analyses, we also generated multivariate rank histograms to thoroughly assess

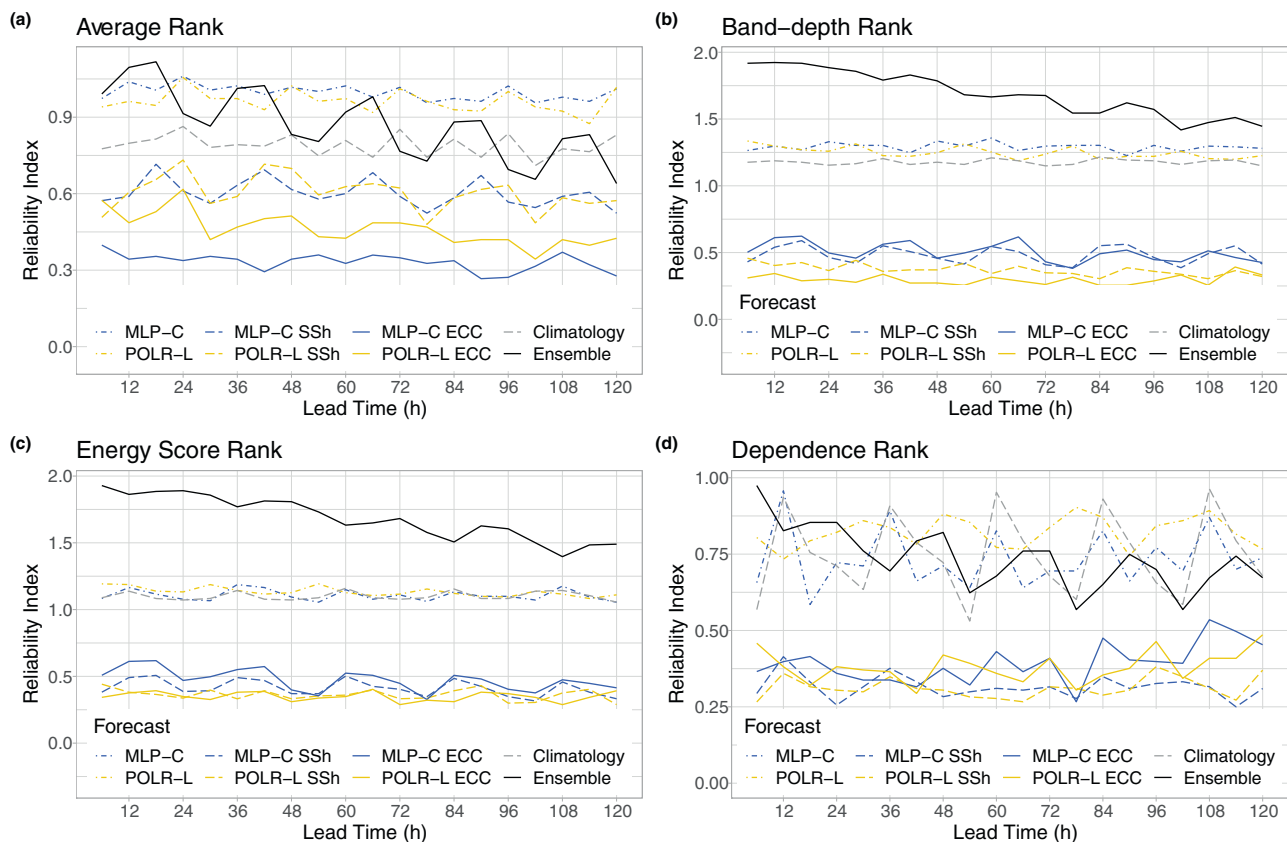


FIGURE 13 Reliability indices of rank histograms of raw, climatological and post-processed multivariate visibility forecasts based on CAMS extended MLP-C and POLR-L models as functions of the lead time.

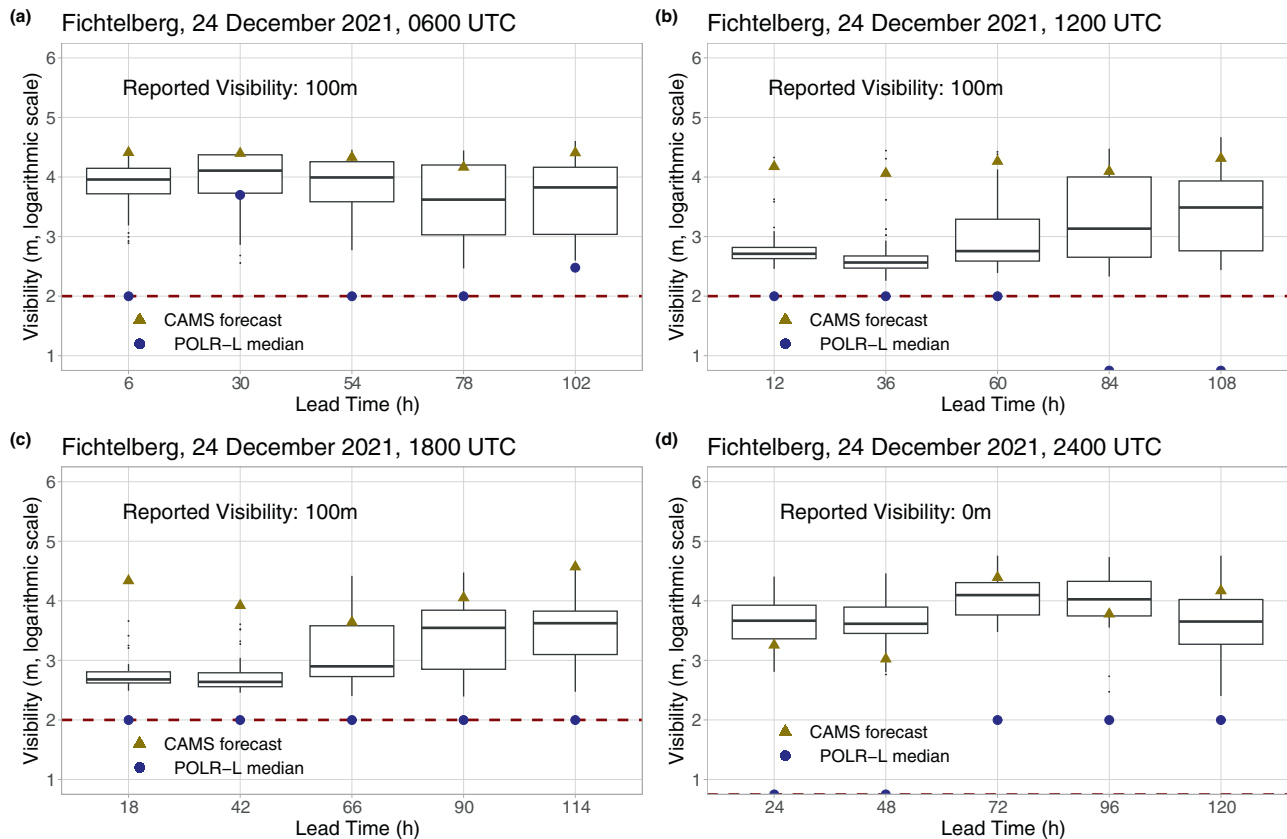


FIGURE 14 Boxplots of ECMWF visibility forecasts (logarithmic scale) for Fichtelberg mountain (Germany) for 0060 UTC (a), 1200 UTC (b), 1800 UTC (c), and 2400 UTC (d) on the Christmas Eve of 2021 together with the corresponding CAMS forecasts and the median forecasts of the CAMS extended POLR-L model. The observed visibility is indicated by the horizontal dashed line. Values on the bottom line of panels (b) and (d) correspond to 0 m visibility.

the differences in calibration of the individual forecasts (see Figure 12). In general, properly calibrated forecasts result in flat histograms, and the deviation from uniformity can be quantified by the reliability index (5). In principle, one can have an instant ranking of the various multivariate predictions with the help of the reliability indices, displayed in Figure 13 (the smaller the better); however, the interpretation of the different rank histograms can vary depending on the functions applied to determine the pre-ranks (see Section 3.4), as these functions may be more sensitive to detecting different deficiencies. Therefore, as argued by Thorarinsdottir et al. (2016) and Allen et al. (2024), using multiple pre-rank functions to evaluate the performance of the forecasts could be beneficial, providing a more comprehensive comparison. In the case of naive MLP-C and POLR-L models, hump-shaped average- and band-depth rank histograms indicate overestimated correlations, which is also supported by the skewed energy score- and dependence rank histograms. To a smaller extent, the same applies to the multivariate climatology. Multivariate calibration results in a substantial improvement, especially

in the case of MLP-C ECC and POLR-L ECC forecasts. According to Figure 13, in terms of the average ranks, these two methods provide by far the lowest reliability indices and are very competitive with the MLP-C SSH and POLR-L SSH approaches in the three other cases, as well.

5 | CONCLUSIONS

The current work is a direct continuation of Baran and Lakatos (2023) where location-specific (univariate) discrete post-processing of visibility was studied. To get a deeper insight into the behavior of this complex and hardly predictable weather quantity, three main questions are investigated. First, we assess whether the state-of-the-art multivariate post-processing approaches which were approved to be successful in the case of more traditional variables such as temperature, wind speed, or precipitation (see Lakatos et al., 2023) can reliably restore spatial dependencies that were lost during separate calibration of visibility ensemble forecasts for each SYNOP

station. Such spatially consistent visibility predictions might have importance, for instance, in route planning related to road transportation or ship navigation. Furthermore, we also investigate whether the use of CAMS visibility predictions as additional covariates in post-processing models can significantly improve the forecast skill. Finally, we study the predictive performance of the post-processed forecasts in the context of aeronautical meteorology.

In the case of the best-performing univariate post-processing approaches (MLP-C and POLR-L) from the methods investigated by Baran and Lakatos (2023), the addition of CAMS predictions to the input features results in a significant 1–5.25% improvement in mean CRPS and 0.27%–1.37% in mean LogS. Note, that for both scores the gain displays a decreasing trend as the lead time increases.

To illustrate the effect of statistical post-processing, consider Christmas Eve of 2021 at the Fichtelberg mountain in Germany when the reported visibility was at most 100 m during the whole day. As depicted in Figure 14, ECMWF forecasts provide good results only for 1200 UTC and 1800 UTC observations and short lead times; CAMS predictions underperform the ECMWF ensemble in the majority of the forecast cases, whereas the median of the CAMS extended POLR-L forecast is very accurate.

This latter approach also provides a powerful tool for predicting the probability of low visibility events, which are of utmost importance in aeronautical applications. Compared to climatology (which outperforms the raw forecast), the average improvement in BS in predicting situations when visibility is below 200, 500, 1000, and 3000 m is between 14.22% and 18.59%.

Joint multivariate post-processing of forecasts for all 30 investigated SYNOP observation stations is performed with the help of the two-step ECC and SSh approaches utilizing the dependence structure of the raw vector ensemble forecasts and historical observations, respectively. In fact, in terms of ES and VS, there are no visible differences in skill between the two dependence template selection methods which substantially outperform the reference independently calibrated (naive) multivariate forecasts. Compared to forecasts based only on the ECMWF predictions, utilizing CAMS results in an additional 0.58–4.11% improvement in energy score for all investigated models. The VS focusing on the correct specification of correlations shows a different picture. In the case of the independent reference methods, the advantage of CAMS extended forecasts is around 5%; however, the corresponding MLP-C forecasts underperform climatology for almost all lead times. In contrast, for the ECC and SSh methods, the VSS of CAMS extended predictions with respect to the models based

only on the ECMWF ensemble is below 2.5%; nevertheless, they outperform climatology by more than 20%. From the competing multivariate predictions, the ECC-corrected MLP-C and POLR-L methods display the best overall performance.

Our case study demonstrated the usefulness of CAMS visibility predictions as additional covariates in the discrete post-processing of visibility ensemble forecasts. Furthermore, it showed the efficiency of the simplest two-step multivariate post-processing methods in capturing spatial dependencies in the case of this particular weather quantity.

Although the current study is based on data from a single limited geographical area and a specific two-year time frame, we believe that our conclusions on the skill of post-processed forecasts are rather general. The results of Baran and Lakatos (2023) confirm that the proposed MLP and POLR approaches also perform well when a different ensemble domain and time period is considered, while the efficiency of the investigated two-step multivariate methods has been proved in several situations involving various weather quantities. In particular, a variant of ECC developed for predicting hourly rainfall is currently in operational use at Météo France (Taillardat & Mestre, 2020).

The presented classification-based univariate post-processing methods are limited to discrete observations, and the more visibility categories are considered, the more training data are required. However, the resolution of the state-of-the-art visibility sensors is 1 m (matching the resolution of the ECMWF and CAMS forecasts), allowing visibility to be treated as a continuous quantity. Thus, a natural direction of further studies is the utilization of CAMS forecasts and/or other visibility-related covariates in continuous post-processing models for visibility ensemble forecasts. Possible parametric candidates are the BMA approach of Chmielecki and Raftery (2011), where the predictive distribution for visibility is a mixture of beta laws, and the more recent censored gamma and censored truncated normal mixture model proposed by Baran and Baran (2024). However, the most straightforward extension of the present work would be the adaptation of the classification and interpolation-based approach of Scheuerer et al. (2020) to visibility forecasts.

AUTHOR CONTRIBUTIONS

Mária Lakatos: Conceptualization; data curation; formal analysis; investigation; software; validation; visualization; writing – original draft; writing – review and editing; methodology. **Sándor Baran:** Conceptualization; funding acquisition; methodology; investigation; validation; visualization; supervision; formal analysis; writing – review and editing.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the ÚNKP-23-3 New National Excellence Program of the Ministry for Culture and Innovation from the source of the National Research, Development and Innovation Fund and of the National Research, Development and Innovation Office under Grant No. K142849. They are also indebted to Martin Leutbecher and Zied Ben Boual-lègue for the inspiring discussions and for providing the ECMWF and CAMS visibility data. Last but not least, the authors thank the three anonymous reviewers for their valuable suggestions improving the manuscript.

FUNDING INFORMATION

ÚNKP-23-3 New National Excellence Program of the Ministry for Culture and Innovation from the source of the National Research, Development and Innovation Fund; National Research, Development and Innovation Office, Grant/Award Number: K142849.

CONFLICT OF INTEREST STATEMENT

None of the authors have a conflict of interest to disclose.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are confidential; however, for research purposes they are available on request from the ECMWF.

ORCID

Sándor Baran  <https://orcid.org/0000-0003-1035-004X>

REFERENCES

- Allen, S., Ziegel, J. & Ginsbourger, D. (2024) Assessing the calibration of multivariate probabilistic forecasts. *Quarterly Journal of the Royal Meteorological Society*, 150, 1315–1335.
- Baran, Á. & Baran, S. (2024) Parametric model for post-processing visibility ensemble forecasts. *Advances in Statistical Climatology, Meteorology and Oceanography*, 10, 105–122.
- Baran, Á., Lerch, S., El Ayari, M. & Baran, S. (2021) Machine learning for total cloud cover prediction. *Neural Computing and Applications*, 33, 2605–2620.
- Baran, S. & Lakatos, M. (2023) Statistical post-processing of visibility ensemble forecasts. *Meteorological Applications*, 30, paper e2157.
- Baran, S. & Möller, A. (2015) Joint probabilistic forecasting of wind speed and temperature using Bayesian model averaging. *Environmetrics*, 26, 120–132.
- Bari, D., Bergot, T. & El Khelifi, M. (2016) Local meteorological and large-scale weather characteristics of fog over the Grand Casablanca region, Morocco. *The Journal of Applied Meteorology and Climatology*, 55, 1731–1745.
- Boudala, F.S., Isaac, G.A., Crawford, R.W. & Reid, J. (2012) Parameterization of runway visual range as a function of visibility: implications for numerical weather prediction models. *Journal of Atmospheric and Oceanic Technology*, 29, 177–191.
- Bremnes, J.B. (2019) Constrained quantile regression splines for ensemble postprocessing. *Monthly Weather Review*, 147, 1769–1780.
- Bremnes, J.B. (2020) Ensemble postprocessing using quantile function regression based on neural networks and Bernstein polynomials. *Monthly Weather Review*, 148, 403–414.
- Chen, J., Janke, T., Steinke, F. & Lerch, S. (2024) Generative machine learning methods for multivariate ensemble postprocessing. *The Annals of Applied Statistics*, 18, 159–183.
- Chmielecki, R.M. & Raftery, A.E. (2011) Probabilistic visibility forecasting using Bayesian model averaging. *Monthly Weather Review*, 139, 1626–1636.
- Clark, M., Gangopadhyay, S., Hay, L., Rajagopalan, B. & Wilby, R. (2004) The Schaake shuffle: a method for reconstructing space-time variability in forecasted precipitation and temperature fields. *Journal of Hydrometeorology*, 5, 243–262.
- Dabernig, M., Mayr, G.J., Messner, J.W. & Zeileis, A. (2017) Spatial ensemble post-processing with standardized anomalies. *Quarterly Journal of the Royal Meteorological Society*, 143, 909–916.
- Delle Monache, L., Hacker, J.P., Zhou, Y., Deng, X. & Stull, R.B. (2006) Probabilistic aspects of meteorological and ozone regional ensemble forecasts. *Journal of Geophysical Research*, 111, paper D24307. Available from: <https://doi.org/10.1029/2005JD006917>
- Dietz, S.J., Kneringer, P., Mayr, G.J. & Zeileis, A. (2019) Low-visibility forecasts for different flight planning horizons using tree-based boosting models. *Advances in Statistical Climatology, Meteorology and Oceanography*, 5, 101–114.
- ECMWF. (2021) *IFS documentation CY47R3 - part IV physical processes*. Reading: ECMWF. Available from: <https://doi.org/10.21957/eyrpir4vj> [Accessed 19th June 2024].
- Friederichs, P. & Hense, A. (2007) Statistical downscaling of extreme precipitation events using censored quantile regression. *Monthly Weather Review*, 135, 2365–2378.
- Gneiting, T. & Raftery, A.E. (2007) Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association*, 102, 359–378.
- Gneiting, T., Raftery, A.E., Westveld, A.H. & Goldman, T. (2005) Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133, 1098–1118.
- Gneiting, T., Stanberry, L.I., Grimit, E.P., Held, L. & Johnson, N.A. (2008) Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *Test*, 17, 211–235.
- Good, I.J. (1952) Rational decisions. *The Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, 14, 107–114.
- Goodfellow, I., Bengio, Y. & Courville, A. (2016) *Deep learning*. Cambridge: MIT Press.
- He, D., Wang, Y., Tang, Y., Kong, D., Yang, J., Zhou, W. et al. (2024) Improvement in the forecasting of low visibility over Guizhou, China, based on a multi-variable deep learning model. *Atmosphere*, 15, paper 752. Available from: <https://doi.org/10.3390/atmos15070752>
- Hemri, S., Haiden, T. & Pappenberger, F. (2016) Discrete postprocessing of total cloud cover ensemble forecasts. *Monthly Weather Review*, 144, 2565–2577.
- Jordan, A., Krüger, F. & Lerch, S. (2019) Evaluating probabilistic forecasts with scoringRules. *Journal of Statistical Software*, 90, 1–37.

- Knüppel, M., Krüger, F. & Pohle, M.-O. (2022) Score-based calibration testing for multivariate forecast distributions. Preprint. <https://doi.org/10.48550/arXiv.2211.16362> [Accessed 19th June 2024].
- Lakatos, M., Lerch, S., Hemri, S. & Baran, S. (2023) Comparison of multivariate post-processing methods using global ECMWF ensemble forecasts. *Quarterly Journal of the Royal Meteorological Society*, 149, 856–877.
- 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 Processes in Geophysics*, 27, 23–34.
- Lerch, S. & Baran, S. (2017) Similarity-based semi-local estimation of EMOS models. *The Journal of the Royal Statistical Society, Series C (Applied Statistics)*, 66, 29–51.
- Lerch, S., Baran, S., Möller, A., Groß, J., Schefzik, R., Hemri, S. et al. (2020) Simulation-based comparison of multivariate ensemble post-processing methods. *Nonlinear Processes in Geophysics*, 27, 349–371.
- López-Pintado, S. & Romo, J. (2009) On the concept of depth for functional data. *Journal of the American Statistical Association*, 104, 718–734.
- McCullagh, P. (1980) Regression model for ordinal data (with discussion). *The Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, 42, 243–268.
- Murphy, A.H. (1973) Hedging and skill scores for probability forecasts. *Journal of Applied Meteorology*, 12, 215–223.
- Penov, N. & Guerova, G. (2023) Sofia airport visibility estimation with two machine-learning techniques. *Remote Sensing*, 15, paper 4799. Available from: <https://doi.org/10.3390/rs15194799>
- Pitkänen, M.R.A., Benedictow, A., Arola, A., Bennouna, Y., Blake, L., Bouarar, I. et al. (2023) Validation report of the CAMS near-real-time global atmospheric composition service: March–May 2023. Copernicus Atmosphere Monitoring Service (CAMS) Report, September 2023. <https://doi.org/10.24380/ac0r-ao9u> [Accessed 19th June 2024].
- Politis, D.N. & Romano, J.P. (1994) The stationary bootstrap. *Journal of the American Statistical Association*, 89, 1303–1313.
- Raftery, A.E., Gneiting, T., Balabdaoui, F. & Polakowski, M. (2005) Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133, 1155–1174.
- Rasp, S. & Lerch, S. (2018) Neural networks for postprocessing ensemble weather forecasts. *Monthly Weather Review*, 146, 3885–3900.
- Roquelaure, S. & Bergot, T. (2009) Contributions from a local ensemble prediction system (LEPS) for improving fog and low cloud forecasts at airports. *Weather and Forecasting*, 24, 39–52.
- Schefzik, R. & Möller, A. (2018) Ensemble postprocessing methods incorporating dependence structures. In: Vannitsem, S., Wilks, D.S. & Messner, J.W. (Eds.) *Statistical Postprocessing of ensemble forecasts*. Amsterdam: Elsevier, pp. 91–125.
- Schefzik, R., Thorarinsdottir, T.L. & Gneiting, T. (2013) Uncertainty quantification in complex simulation models using ensemble copula coupling. *Statistical Science*, 28, 616–640.
- Scheuerer, M. & Hamill, T.M. (2015) Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Monthly Weather Review*, 143, 1321–1334.
- Scheuerer, M., Switanek, M.B., Worsnop, R.P. & Hamill, T.M. (2020) Using artificial neural networks for generating probabilistic subseasonal precipitation forecasts over California. *Monthly Weather Review*, 148, 3489–3506.
- Schultz, B. & Lerch, S. (2022) Machine learning methods for post-processing ensemble forecasts of wind gusts: a systematic comparison. *Monthly Weather Review*, 150, 235–257.
- Shankar, A. & Sahana, B.C. (2023) Early warning of low visibility using the ensembling of machine learning approaches for aviation services at Jay Prakash Narayan International (JPNI) Airport Patna. *SN Applied Sciences*, 5, paper 132. Available from: <https://doi.org/10.1007/s42452-023-05350-7>
- Taillardat, M. & Mestre, O. (2020) From research to applications—examples of operational ensemble post-processing in France using machine learning. *Nonlinear Processes in Geophysics*, 27, 329–347.
- Taillardat, M., Mestre, O., Zamo, M. & Naveau, P. (2016) Calibrated ensemble forecasts using quantile regression forests and ensemble model output statistics. *Monthly Weather Review*, 144, 2375–2393.
- Thorarinsdottir, T.L. & Gneiting, T. (2010) Probabilistic forecasts of wind speed: ensemble model output statistics by using heteroscedastic censored regression. *The Journal of the Royal Statistical Society, Series A (Statistics in Society)*, 173, 371–388.
- Thorarinsdottir, T.L., Scheuerer, M. & Heinz, C. (2016) Assessing the calibration of high-dimensional ensemble forecasts using rank histograms. *Journal of Computational and Graphical Statistics*, 25, 105–122.
- Vannitsem, S., Bremnes, J.B., Demaeyer, J., Evans, G.R., Flowerdew, J., Hemri, S. et al. (2021) Statistical postprocessing for weather forecasts—review, challenges and avenues in a big data world. *Bulletin of the American Meteorological Society*, 102, E681–E699.
- Wilks, D.S. (2019) *Statistical methods in the atmospheric sciences*, 4th edition. Amsterdam: Elsevier.
- World Meteorological Organization. (1992) *International meteorological vocabulary (WMO-No. 182)*. Geneva: WMO.
- World Meteorological Organization. (2018) *Guide to instruments and methods of observation. Volume I—Measurement of meteorological variables (WMO-No. 8)*. Geneva: WMO.
- World Meteorological Organization. (2022) Aerodrome reports and forecasts: a users. In: *Handbook to the codes (WMO-No. 782)*. Geneva: WMO.

How to cite this article: Lakatos, M., & Baran, S. (2024). Enhancing multivariate post-processed visibility predictions utilizing Copernicus Atmosphere Monitoring Service forecasts. *Meteorological Applications*, 31(6), e70015. <https://doi.org/10.1002/met.70015>