

GENERATIVE MACHINE LEARNING METHODS FOR MULTIVARIATE ENSEMBLE POSTPROCESSING

BY JIEYU CHEN^{1,a} , TIM JANKE^{2,c} , FLORIAN STEINKE^{2,d}  AND SEBASTIAN LERCH^{1,b} 

¹*Institute of Economics, Karlsruhe Institute of Technology, jieyu.chen@kit.edu, sebastian.lerch@kit.edu*

²*Energy Information Networks & Systems, Technical University of Darmstadt, tim.janke@gmx.net, florian.steinke@eins.tu-darmstadt.de*

Ensemble weather forecasts based on multiple runs of numerical weather prediction models typically show systematic errors and require postprocessing to obtain reliable forecasts. Accurately modeling multivariate dependencies is crucial in many practical applications, and various approaches to multivariate postprocessing have been proposed where ensemble predictions are first postprocessed separately in each margin and multivariate dependencies are then restored via copulas. These two-step methods share common key limitations, in particular, the difficulty to include additional predictors in modeling the dependencies. We propose a novel multivariate postprocessing method based on generative machine learning to address these challenges. In this new class of nonparametric data-driven distributional regression models, samples from the multivariate forecast distribution are directly obtained as output of a generative neural network. The generative model is trained by optimizing a proper scoring rule, which measures the discrepancy between the generated and observed data, conditional on exogenous input variables. Our method does not require parametric assumptions on univariate distributions or multivariate dependencies and allows for incorporating arbitrary predictors. In two case studies on multivariate temperature and wind speed forecasting at weather stations over Germany, our generative model shows significant improvements over state-of-the-art methods and particularly improves the representation of spatial dependencies.

1. Introduction. Most weather forecasts today are based on ensemble simulations of numerical weather prediction (NWP) models. Despite substantial improvements over the past decades (Bauer, Thorpe and Brunet (2015)), these ensemble predictions continue to exhibit systematic errors, such as biases, and typically fail to correctly quantify forecast uncertainty. These systematic errors can be corrected by statistical postprocessing, the application of which has become standard practice in research and operations. Over the past years, a focal point of research interest has been the use of modern machine learning (ML) methods for postprocessing, where random forest or neural network models enable the incorporation of arbitrary input predictors and have demonstrated superior forecast performance (Rasp and Lerch (2018), Taillardat et al. (2016)). For a general overview of recent developments, we refer to Vannitsem et al. (2021) and Haupt et al. (2021).

While much of the research interest in postprocessing has been focused on univariate methods, many practical applications require accurate models of spatial, temporal, or inter-variable dependencies (Schefzik, Thorarinsdottir and Gneiting (2013)). Key examples include energy forecasting (Pinson and Messner (2018), Worsnop et al. (2018)), air traffic management (Chaloulos and Lygeros (2007)) and hydrological applications (Scheuerer et al. (2017)). While the spatial, temporal, or intervariable dependencies are present in the raw ensemble

Received September 2022; revised May 2023.

Key words and phrases. Generative machine learning, ensemble postprocessing, multivariate postprocessing, probabilistic forecasting, weather forecasting.

predictions from the NWP model, they are lost if univariate postprocessing methods are applied separately in each margin (i.e., at each location, time step and for each target variable), which corresponds to implicitly assuming independence across space, time and variables.

Over the past decade, various multivariate postprocessing methods have been proposed; see [Schefzik and Möller \(2018\)](#) and [Vannitsem et al. \(2021\)](#) for general overviews. The vast majority of multivariate postprocessing methods follows a two-step strategy.¹ In the first step, univariate postprocessing methods are applied independently in all dimensions to obtain calibrated marginal probability distributions. Samples are then generated from these marginal predictive distributions obtained after postprocessing. In the second step, the univariate sample values are rearranged, according to a specific multivariate dependence template, with the aim of restoring the multivariate dependencies that are lost in the first step. From a mathematical perspective, this can be interpreted as the application of a (parametric or nonparametric) copula, that is, a multivariate cumulative distribution function (CDF) with standard uniform marginal distributions ([Nelsen \(2006\)](#)).

In most popular copula-based approaches to multivariate postprocessing, the copula C is chosen to be either the parametric Gaussian copula (in the Gaussian copula approach, GCA; [Möller, Lenkoski and Thorarinsdottir \(2013\)](#), [Pinson and Girard \(2012\)](#)) or a nonparametric empirical copula induced by a prespecified dependence template based on the raw ensemble predictions (in the ensemble copula coupling (ECC; [Schefzik, Thorarinsdottir and Gneiting \(2013\)](#)) approach) or past observations (in the Schaake shuffle approach, [Clark et al. \(2004\)](#)). More advanced variants of these methods, which aim to better incorporate structure in forecast error autocorrelations ([Bouallègue et al. \(2016\)](#)) or optimize the selection of the dependence template based on similarity ([Schefzik \(2016\)](#), [Scheuerer et al. \(2017\)](#)), have been proposed over the past years. Several comparative studies of multivariate postprocessing methods, based on simulated and real-world data, are available ([Lakatos et al. \(2023\)](#), [Lerch et al. \(2020\)](#), [Perrone, Schicker and Lang \(2020\)](#), [Wilks \(2015\)](#)). Overall, findings from these studies indicate that there is no consistently best approach across all settings and that the observed differences in predictive performance depend on the misspecifications of the raw ensemble predictions, but tend to be minor and not statistically significant.

All these state-of-the-art two-step methods for multivariate postprocessing share common key limitations. Perhaps most importantly, there is no straightforward way to include additional predictors beyond ensemble forecasts or past observations of the target variable in the second step of imposing the dependence template onto the samples from the univariate post-processed forecast distributions. Incorporating additional predictors has proven to be a key aspect in the substantial improvements in predictive performance observed for ML-based univariate postprocessing models ([Rasp and Lerch \(2018\)](#), [Taillardat et al. \(2016\)](#), [Vannitsem et al. \(2021\)](#)), and it seems reasonable to expect similar beneficial effects for modeling multivariate dependencies. Furthermore, the number of samples that can be obtained from the multivariate forecast distribution in the two-step postprocessing methods is limited by the number of ensemble predictions (in the ECC approach) or the number of suitable past observations (in case of the Schaake shuffle), which might be disadvantageous for reliably representing and predicting multivariate extreme events.

¹There exist examples of directly fitting specific multivariate probability distributions, which are mostly used in low-dimensional settings or if a specific structure can be chosen for the application at hand, with examples focusing on spatial ([Feldmann, Scheuerer and Thorarinsdottir \(2015\)](#)), temporal ([Muschinski et al. \(2022\)](#)) and intervariable ([Baran and Möller \(2015\)](#), [Lang et al. \(2019\)](#), [Schuhen, Thorarinsdottir and Gneiting \(2012\)](#)) dependencies being available. In addition, there are alternative approaches, such as member-by-member postprocessing ([Van Schaeybroeck and Vannitsem \(2015\)](#)), which inherently preserve dependencies present in the raw ensemble predictions.

To overcome these challenges, we propose a novel nonparametric multivariate postprocessing method based on generative ML approaches. In this new class of data-driven multivariate distributional regression models, samples from the multivariate forecast distribution are directly obtained as output of a generative deep neural network which allows for incorporating arbitrary input predictors including NWP-based ensemble predictions and additional exogenous variables. Our generative model circumvents the two-step structure of the state-of-the-art multivariate postprocessing approaches and aims to simultaneously correct systematic errors in the marginal distributions and the multivariate dependence structure without requiring parametric assumptions.

There has recently been an increase in research activity on generative machine learning methods for weather modeling, in particular, in the context of downscaling precipitation forecasts (Harris et al. (2022), Hess et al. (2022), Leinonen, Nerini and Berne (2021), Price and Rasp (2022)) and nowcasting (Ravuri et al. (2021)), where several approaches based on generative adversarial networks (GANs) have been proposed. While GANs are particularly useful for generating realistic images and have shown some success in a postprocessing application to total cloud cover forecasts (Dai and Hemri (2021)), their training is often challenging (Gui et al. (2021)). We consider a conceptually simpler class of scoring rule-based generative models (Li, Swersky and Zemel (2015), Dziugaite, Roy and Ghahramani (2015)), extending recent work in probabilistic model averaging in energy forecasting (Janke and Steinke (2020)).² In our approach a conditional generative model, based on a neural network, is trained by optimizing a suitable multivariate proper scoring rule, which measures the discrepancy between the generated and true data, and replaces the GAN discriminator. By conditioning the generative process on exogenous input variables, we enable the generative model to incorporate arbitrary input predictors beyond random noise only and to flexibly learn nonlinear relations to both the univariate forecast distributions and the multivariate dependence structure.

Using two case studies on spatial dependencies of temperature and wind speed forecasts at weather stations over Germany, we compare our conditional generative model with state-of-the-art two-step approaches to multivariate postprocessing based on ECC and GCA. For the univariate postprocessing step, we consider models only based on ensemble predictions of the target variable as well as state-of-the-art neural network models with additional predictors. This will allow for specifically investigating the benefits of incorporating additional information in the marginal distribution or the full multivariate forecast.

The remainder of the paper is organized as follows. Section 2 introduces the datasets, and Section 3 reviews the standard two-step postprocessing methods. Our conditional generative model is described in Section 4, followed by the main results presented in Section 5. Section 6 concludes with a discussion. Python and R code with implementations of all methods is available in the Supplementary Material (Chen et al. (2024)) and online (<https://github.com/jieyu97/mvpp>).

2. Data. Our study focuses on forecasts of 2-m temperature and 10-m wind speed with a forecast lead time of 48 hours. The dataset for temperature³ is based on the one used in Rasp and Lerch (2018), while the dataset for wind speed⁴ was compiled for this study. Both datasets are based on forecasts from the 50-member ensemble of the European Center of Medium-range Weather Forecasts (ECMWF) initialized at 00 UTC every day, which were obtained from the THORPEX Interactive Grand Global Ensemble (TIGGE) database (Bougeault, Toth et al. (2010)) on a $0.5^\circ \times 0.5^\circ$ grid over Europe. Following the procedure

²For theoretical results on generative models based on scoring rule optimization, see Pacchiardi et al. (2021).

³The dataset is available from <https://doi.org/10.6084/m9.figshare.19453580>.

⁴The dataset is available from <https://doi.org/10.6084/m9.figshare.19453622>.




FIG. 1. Locations of weather stations with (a) temperature and (b) wind speed observations.

outlined in Rasp and Lerch (2018), the ensemble forecasts of all meteorological predictor variables are interpolated to weather station locations over Germany. Stations with larger fractions of missing data and with altitudes above 1000 m are omitted to avoid outliers, due to substantially different topographical properties when considering spatial dependencies. This results in a total of 419 stations in the temperature dataset and 198 stations in the wind speed dataset; the locations of which are shown in Figure 1. Corresponding observations were obtained from the Climate Data Center⁵ of the German weather service.

In addition to ensemble forecasts of the target variables (temperature and wind speed, respectively), ensemble forecasts of several auxiliary predictor variables, based on the selection in Rasp and Lerch (2018), are available. Ensemble forecasts of all meteorological predictor variables are reduced to their mean and standard deviation. For the wind speed dataset, we also explicitly compute the wind speeds at different pressure levels from the corresponding wind components. In addition, we use a sine-transformed value of the day of the year⁶ and relevant information about the station coordinates, altitudes and orography (altitude of the model grid point) as additional input predictors for the postprocessing models. Table 1 provides an overview of all available predictors in both datasets.

For both datasets a total of 10 years of daily forecast and observation data from 2007–2016 are available. Following Rasp and Lerch (2018), we use data from 2007–2014 as training set and 2015 as validation set for choosing hyperparameters and model specifications. Data from 2016 serves as out-of-sample test dataset and is not used during model training.

3. Benchmark methods for multivariate ensemble postprocessing. This section introduces state-of-the-art approaches to multivariate postprocessing, which serve as benchmark methods for our generative machine learning method that will be introduced in Section 4. We focus on the two-step approaches based on a combination of univariate postprocessing models and copulas. In a first step, univariate postprocessing is applied to ensemble forecasts for each margin (i.e., location, forecast horizon or target variable), and in a second step, the multivariate dependence structure is imposed upon the univariately postprocessed forecast using a suitable copula. Sklar’s theorem (Sklar (1959)) provides the theoretical underpinning in that a multivariate CDF H (our target) can be decomposed into a copula

⁵https://www.dwd.de/DE/klimaumwelt/cdc/cdc_node.html.

⁶The numerical value of the day of the year ranging from 1 to 366, t is transformed via $\text{doy} = \sin\left(\frac{t-105}{366} \cdot 2\pi\right)$.

TABLE 1

Available predictors for temperature and wind speed postprocessing. The variables ws , ws_pl850 and ws_pl500 are derived from the corresponding U and V wind components and are not included in the temperature dataset. For wind speed forecasts, the orography ($orog$) information is missing

Variable	Description
<i>Meteorological variables</i>	
t2m	2-m temperature
d2m	2-m dewpoint temperature
cape	Convective available potential energy
sp	Surface pressure
tcc	Total cloud cover
q_pl850	Specific humidity at 850 hPa
u_pl850	U component of wind at 850 hPa
v_pl850	V component of wind at 850 hPa
ws_pl850	Wind speed at 850 hPa
u_pl500	U component of wind at 500 hPa
v_pl500	V component of wind at 500 hPa
gh_pl500	Geopotential height at 500 hPa
ws_pl500	Wind speed at 500 hPa
u10	10-m U component of wind
v10	10-m V component of wind
ws	10-m wind speed
sshf	Sensible heat flux
slhf	Latent heat flux
ssr	Shortwave radiation flux
str	Longwave radiation flux
<i>Other predictors</i>	
lon	Longitude of station
lat	Latitude of station
alt	Altitude of station
orog	Altitude of model grid point
doy	Sine-transformed value of the day of the year

function C modeling the dependence structures (this is what needs to be specified) and its marginal univariate CDFs F_1, \dots, F_D (here obtained via univariate postprocessing) via

$$H(z_1, \dots, z_D) = C(F_1(z_1), \dots, F_D(z_D))$$

for $z_1, \dots, z_D \in \mathbb{R}$.

In the following we first introduce methods for univariate postprocessing and then discuss the use of copula functions for restoring multivariate dependencies. Our notation follows that of [Lerch et al. \(2020\)](#), and we denote the unprocessed D -dimensional ensemble forecasts of a weather variable with M members by $\mathbf{X}_1, \dots, \mathbf{X}_M \in \mathbb{R}^D$, where $\mathbf{X}_m = (X_m^{(1)}, \dots, X_m^{(D)})$ for $m = 1, \dots, M$. The corresponding observation is $\mathbf{y} = (y^{(1)}, \dots, y^{(D)}) \in \mathbb{R}^D$, and we use a generic index $d = 1, \dots, D$ to summarize the margins (in our case the locations when modeling spatial dependencies).

3.1. Univariate postprocessing. Over the past years, a large variety of univariate postprocessing methods have been proposed. In particular, the development of modern methods from ML for postprocessing has been a focus of recent research interest. We refer to [Vannitsem et al. \(2021\)](#) for a general overview and to [Schulz and Lerch \(2022a\)](#) for a recent comparison in the context of wind gust forecasting. We restrict our attention to postprocessing methods within the parametric distributional regression framework proposed by [Gneiting et al. \(2005\)](#).

To specifically investigate the effect of improved marginal predictions, we consider a simple ensemble model output statistics approach (Gneiting et al. (2005)) and a state-of-the-art method based on neural networks (Rasp and Lerch (2018)), which allows for incorporating additional predictor information and flexibly model nonlinear relations to forecast distribution parameters.

Within the distributional regression framework, the conditional probability distribution of the (univariate) variable of interest y , given ensemble predictions X_1, \dots, X_M of the target variable, is modeled by a parametric distribution, F_θ , the parameters of which depend on the ensemble predictions via a link function g , that is,

$$\theta = g(X_1, \dots, X_M).$$

Note that, within the current subsection, we suppress the index of the dimension to simplify notation.

3.1.1. *Ensemble model output statistics (EMOS)*. Following Gneiting et al. (2005), the standard EMOS model for temperature (t2m) assumes a Gaussian predictive distribution \mathcal{N} with mean μ and standard deviation σ , that is,

$$y|X_1^{\text{t2m}}, \dots, X_M^{\text{t2m}} \sim \mathcal{N}(\mu, \sigma),$$

and uses ensemble forecasts of the target variable, $X_1^{\text{t2m}}, \dots, X_M^{\text{t2m}}$, as sole predictors. The distribution parameters are linked to the ensemble mean and standard deviation (sd) via

$$(1) \quad \mu = a_0 + a_1 \cdot \text{mean}(X_1^{\text{t2m}}, \dots, X_M^{\text{t2m}}) \quad \text{and} \quad \sigma = b_0 + b_1 \cdot \text{sd}(X_1^{\text{t2m}}, \dots, X_M^{\text{t2m}}).$$

The EMOS coefficients a_0, a_1, b_0, b_1 are determined by minimizing the continuous ranked probability score (CRPS; Matheson and Winkler (1976), see Section 5.1 for details) on the training dataset. We here implement local EMOS models by estimating separate sets of coefficients for different stations, which makes the models locally adaptive and typically leads to better performance compared with a global model that is jointly estimated for all stations, provided that a sufficient amount of training data is available (Lerch and Baran (2017)). While rolling training windows consisting of only the most recent days have often been used for the estimation of EMOS models, we use a static training period, given by the entire training dataset, following common practice in the operational use of postprocessing models. Furthermore, several studies suggest that the benefits of using long archives of training data often outweigh potential changes in the underlying NWP model or the meteorological conditions (e.g., Lang et al. (2020)).

For wind speed we proceed analogously but follow Thorarinsdottir and Gneiting (2010) and utilize a Gaussian distribution left-truncated at 0. This ensures that no probability mass is assigned to negative wind speed values. Forecasts of wind speed (ws) are used as sole input predictors, that is,

$$y|X_1^{\text{ws}}, \dots, X_M^{\text{ws}} \sim \mathcal{N}_{[0, \infty)}(\mu, \sigma),$$

where $\mathcal{N}_{[0, \infty)}(\mu, \sigma)$ denotes a truncated Gaussian distribution with cumulative distribution function

$$F(z) = \Phi\left(\frac{\mu}{\sigma}\right)^{-1} \Phi\left(\frac{z - \mu}{\sigma}\right)$$

for $z > 0$ and 0 otherwise. By contrast to the situation for temperature, the choice of a parametric distribution is less clear for wind speed, and a large variety of distributions have been considered (Baran and Lerch (2015), Baran, Szokol and Szabó (2021), Lerch and Thorarinsdottir (2013), Pantillon et al. (2018), Scheuerer and Möller (2015)), including weighted

mixtures of distributions (Baran and Lerch (2016), Sloughter, Gneiting and Raftery (2010)). However, the differences across parametric distributions are usually only minor and are unlikely to effect our results and conclusions. To link the parameters of the truncated Gaussian distribution to the ensemble predictions, we proceed as in (1), that is,

$$(2) \quad \mu = a_0 + a_1 \cdot \text{mean}(X_1^{\text{ws}}, \dots, X_M^{\text{ws}}) \quad \text{and} \quad \sigma = b_0 + b_1 \cdot \text{sd}(X_1^{\text{ws}}, \dots, X_M^{\text{ws}}).$$

3.1.2. *Neural network methods for univariate postprocessing.* Many standard approaches to postprocessing, including the EMOS method introduced above, share a common limitation in that incorporating additional predictors beyond forecasts of the target variable is challenging. To do so, it would be necessary to manually specify the exact functional form of the dependencies between the distribution parameters and all available input predictors in equations (1) and (2). During the past years, a variety of ML methods have been developed to address this issue (Vannitsem et al. (2021)). Rasp and Lerch (2018) propose a neural network (NN) approach, where the distribution parameters are obtained as the output of a NN, which allows for learning arbitrary nonlinear relations between input predictors and distribution parameters in an automated, data-driven manner. We refer to this approach as the distributional regression network (DRN).

In our implementations we follow Rasp and Lerch (2018) and use a NN with one hidden layer. All available predictors, except for the date information, are normalized to the range [0, 1] using a min-max scaler and are then used as inputs to the NN, which returns the distribution parameters μ and σ as outputs. A single NN model is estimated jointly for all stations, using the CRPS as a custom loss function. The model predictions are made locally adaptive by the use of embeddings of the station identifiers, a technique that was originally proposed in natural language processing (Pennington, Socher and Manning (2014)). Our model architecture and implementation choices directly follow those of Rasp and Lerch (2018) for temperature, where we employ a Gaussian predictive distribution. For wind speed we use a truncated Gaussian predictive distribution, as in the corresponding EMOS model, and apply a softplus activation,

$$\text{softplus}(z) = \log(1 + \exp(z)),$$

to the output layer to ensure positivity of the distribution parameters, which helps to avoid numerical issues.

The results, presented in Rasp and Lerch (2018), indicate that the DRN approach to postprocessing leads to substantial improvements over state-of-the-art benchmark methods, and subsequent research has generalized the methodology to other target variables and characterizations of the forecast distribution (Bremnes (2020), Chapman et al. (2022), Scheuerer et al. (2020), Schulz and Lerch (2022a)). For our purposes here, it will be interesting to investigate whether the improvements in the univariate predictive performance directly extends to improved multivariate predictions when coupling the univariate DRN models with the re-ordering techniques introduced below.

3.2. *Multivariate extensions using copulas.* Univariate postprocessing methods are intended to correct systematic errors in the marginal distributions. However, multivariate dependencies are lost when univariate postprocessing is applied separately for each margin (e.g., each station in our application) and need to be restored. A variety of methods for restoring multivariate dependencies via copula functions have been proposed over the past years. We here limit our discussion to the popular ensemble copula coupling and Gaussian copula approach, and refer to Schefzik, Thorarinsdottir and Gneiting (2013), Wilks (2015) and Lerch et al. (2020) for overviews and comparisons. In our descriptions we follow Lerch et al. (2020) and refer to their Section 2 for further mathematical details and references.

3.2.1. *Ensemble copula coupling.* Given univariately postprocessed marginal distributions, a sample of the same size as the raw ensemble, M , is drawn from each predictive marginal distribution. While several sampling schemes have been proposed (Hu et al. (2016), Schefzik, Thorarinsdottir and Gneiting (2013)), we only consider the use of equidistant quantiles at levels $\frac{1}{M+1}, \dots, \frac{M}{M+1}$. Ensemble copula coupling (ECC) is based on the assumption that the ensemble forecasts are informative about the true multivariate dependence structure, and it makes use of the rank order structure of the raw ensemble member forecasts to rearrange the sampled values, with possible ties resolved at random. This can be interpreted as a nonparametric, empirical copula approach, which we refer to as ECC-Q.

A widely used alternative nonparametric approach is the Schaake shuffle (Clark et al. (2004)), which proceeds as ECC, but reorders the sampled quantiles based on the rank order structure of past observations instead of the ensemble forecasts. As noted in the Introduction, comparative studies have often found similar predictive performances between the Schaake shuffle and ECC (e.g., Lakatos et al. (2023)). For the datasets at hand, initial tests indicated a slightly worse performance, compared to ECC-Q (not shown), and we thus only use ECC-Q as a nonparametric benchmark approach to retain focus.

3.2.2. *Gaussian copula approach.* In contrast to ECC-Q, the Gaussian copula approach (GCA; Möller, Lenkoski and Thorarinsdottir (2013), Pinson and Girard (2012)) is based on a parametric Gaussian copula. In a first step of the application of GCA, a set of past observations is transformed into latent standard Gaussian observations via

$$\tilde{y}^{(d)} = \Phi^{-1}(F_{\theta}^{(d)}(y^{(d)}))$$

for all dimensions $d = 1, \dots, D$, where $F_{\theta}^{(d)}$ denotes the corresponding forecast distributions obtained via univariate postprocessing. In a next step, multivariate random samples $\mathbf{Z}_1, \dots, \mathbf{Z}_M$ are randomly drawn from a D -dimensional Gaussian distribution $\mathcal{N}^{(D)}(\mathbf{0}, \Sigma)$ with a mean vector of 0 and an empirical covariance matrix Σ based on the observations transformed into a latent Gaussian space in the first step. The final postprocessed GCA forecasts are then obtained via

$$X_m^{\text{GCA}(d)} = (F_{\theta}^{(d)})^{-1}(\Phi(Z_m^{(d)}))$$

for $m = 1, \dots, M$ and $d = 1, \dots, D$.

In addition to the assumption of a parametric copula, the main difference of GCA to ECC-Q is given by the use of past observations to determine the dependence template. While the number of GCA ensemble members is not limited by the size of the raw ensemble, we here only consider $M = 50$ to ensure comparability across methods.

To summarize, in the following we will consider four two-step approaches based on the available combinations of methods for univariate postprocessing (EMOS, DRN) and copula-based modeling of multivariate dependencies (ECC-Q, GCA) for both target variables (temperature and wind speed), which will serve as benchmarks for our generative ML approach. These benchmark methods will be abbreviated by EMOS+ECC, EMOS+GCA, DRN+ECC and DRN+GCA, respectively.

4. Generative models for multivariate distributional regression. We propose a novel approach to multivariate ensemble postprocessing using generative machine learning techniques. Moving beyond the previous two-step strategy of separately modeling marginal distributions and multivariate dependence structure, this new class of data-driven multivariate distributional regression models allows for obtaining multivariate probabilistic forecasts directly as output of a NN. The proposed generative models provide a nonparametric way to generate postprocessed multivariate samples without any distributional assumptions on the

marginal distributions or the multivariate dependencies. By allowing for incorporating additional predictors beyond forecasts of the target variable and for generating an arbitrary number of samples, the proposed generative models further address key limitations of state-of-the-art two-step approaches to multivariate postprocessing.

4.1. *Deep generative models.* Implicit generative models aim to provide a representation of the probability distribution of a target variable by defining a stochastic procedure to generate samples $Y_1, \dots, Y_{n_{\text{out}}}$ from the distribution of interest (Mohamed and Lakshminarayanan (2016)). The only input to the generative model

$$(3) \quad Y_i = g_{\theta}(Z_i)$$

are samples $Z_1, \dots, Z_{n_{\text{out}}}$ from a simple base distribution, for example, a standard multivariate normal distribution $Z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The learnable map g is typically parameterized by a deep NN with parameters θ .

As implicit generative models do not provide a tractable density of the target distribution, classic parameter estimation procedures, such as maximum likelihood estimation, are infeasible. Generative adversarial networks (Goodfellow et al. (2014)) sidestep this problem by specifying an additional classification model, the discriminator, which is trained to discriminate between the true data and the generated samples. The generator model is then trained to maximize the misclassification rate of the discriminator. In the context of ensemble postprocessing, Dai and Hemri (2021) propose a GAN-based model for generating spatially coherent maps of total cloud cover forecasts.

While impressive results have been achieved by GANs, in particular in the context of image processing and computer vision, the adversarial training process is often considered to be complex and unstable (Gui et al. (2021)). Therefore, several alternative approaches for training generative models have been developed. From a statistical perspective, generative moment matching networks (GMMNs; Li, Swersky and Zemel (2015), Dziugaite, Roy and Ghahramani (2015)) are a particularly interesting class of generative models. GMMNs replace the discriminator with the maximum mean discrepancy (MMD; Gretton et al. (2012)), a kernel-based two-sample test statistic which can be used to measure distances on the space of probability distributions. The training objective is then to minimize a Monte Carlo estimate of the MMD. However, GMMNs only approximate the unconditional data distribution $P(Y)$ while we are interested in the conditional distribution $P(Y|X)$; that is, we aim for a conditional generative model of the form

$$(4) \quad Y_i = g_{\theta}(X, Z_i).$$

Our approach builds on recent work on probabilistic model averaging in energy forecasting (Janke and Steinke (2020)) and uses the energy score (ES; Gneiting and Raftery (2007a), see Section 5.1 for details) as a loss function to train a conditional generative model. The ES is a multivariate strictly proper scoring rule, which is derived from the energy distance (Székely and Rizzo (2013)), which in turn is a special case of the MMD (Sejdinovic et al. (2013)). Training a conditional generative model, based on ES optimization, allows for generating multivariate postprocessed forecasts that simultaneously correct systematic biases and dispersion errors as well as systematic errors in the multivariate dependence structure.

4.2. *Notation.* We now introduce the notation that will be used for describing the model architecture and training process. As input predictors at every weather station location $d = 1, \dots, D$, we have NWP forecasts of K different weather variables available (see Table 1), each in the form of an ensemble of size M ; that is, $X_{k,m}^{(d)}$ is the value of the m th ensemble member for variable a_k at location d with $m = 1, \dots, M$, $k = 1, \dots, K$ and $d = 1, \dots, D$.

The ensemble weather forecasts are reduced to the corresponding ensemble mean $\mu(\mathbf{X}_k^{(d)}) = \frac{1}{M} \sum_m X_{k,m}^{(d)}$ and standard deviation $\sigma(\mathbf{X}_k^{(d)})$. In the following we will use

$$\bar{\mathbf{X}}^{(d)} = [\mu(\mathbf{X}_1^{(d)}), \dots, \mu(\mathbf{X}_K^{(d)})]^T \quad \text{and} \quad \mathbf{s}^{(d)} = [\sigma(\mathbf{X}_1^{(d)}), \dots, \sigma(\mathbf{X}_K^{(d)})]^T$$

to denote vectors of size K , which contain the mean ensemble predictions of all variables at location d and the corresponding standard deviations, respectively. As additional static input predictors, we use the vector $\mathbf{loc}^{(d)} = [\text{lat}^{(d)}, \text{lon}^{(d)}, \text{alt}^{(d)}, \text{orog}^{(d)}]^T$ with location-specific information as well as the (scalar) sine-encoded day of the year (doy). We will denote a single sample from the D_{latent} -dimensional noise distribution by $\mathbf{z}_i = [z_1, \dots, z_{D_{\text{latent}}}]^T \sim \mathcal{N}^{D_{\text{latent}}}(\mathbf{0}, \mathbf{I})$ and denote a single final output sample from the D -dimensional target distribution by $\hat{\mathbf{y}}_i, i = 1, \dots, n_{\text{out}}$.

4.3. Model architecture and training. A schematic overview of our conditional generative model (CGM) is provided in Figure 2. The same basic model structure is used for both temperature and wind speed prediction, and we will highlight relevant differences in the following. The final output of the model is given by D -dimensional multivariate samples drawn from the postprocessed joint distribution, $\{\hat{\mathbf{y}}_i, i = 1, \dots, n_{\text{out}}\}$, which are composed of a mean component $\mathbf{y}^{\text{mean}} \in \mathbb{R}^D$ and a noise component $\mathbf{y}_i^{\text{noise}} \in \mathbb{R}^D$ that depends on the sample $\mathbf{z}_i \in \mathbb{R}^{D_{\text{latent}}}$ from the latent noise distribution for $i = 1, \dots, n_{\text{out}}$.

Our CGM architecture allows for incorporating arbitrary input predictors and for specifically tailoring the model structure to incorporate relevant exogenous information in the different components of the target distribution by separating the mean and noise component. To efficiently propagate the uncertainty inherent to the NWP forecasts, we dynamically reparametrize the latent distributions of the generative model conditionally on the standard deviations of the NWP ensemble forecasts. The overall model consists of three modules with different sets of inputs, which will be described in the following.

The first part of the model aims to learn the multivariate mean component \mathbf{y}^{mean} of the forecast distribution and can be considered as a multivariate deterministic bias-correction step of the mean ensemble forecasts, conditional on the available additional predictors. This

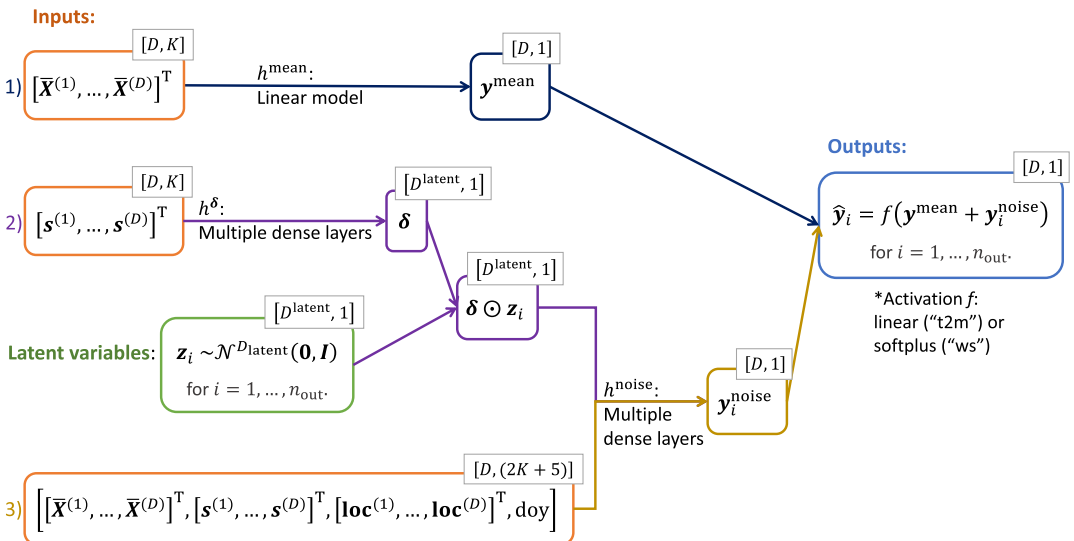


FIG. 2. Schematic illustration of the conditional generative model. The dimensions of the tensors at each step are indicated in the small box.

mean module thus maps mean ensemble forecasts of all weather input variables to \mathbf{y}^{mean} , that is,

$$\mathbf{y}^{\text{mean}} = h^{\text{mean}}(\bar{\mathbf{X}}^{(1)}, \dots, \bar{\mathbf{X}}^{(D)}).$$

We utilize a linear model without hidden layers for h^{mean} ; that is, we have for each dimension $d = 1, \dots, D$

$$y^{\text{mean}(d)} = w_{0,d} + \sum_{k=1}^K w_{k,d} \cdot \mu(\mathbf{X}_k^{(d)}).$$

The remaining two parts of the model aim to learn the noise component $\mathbf{y}^{\text{noise}}$ of the multivariate forecast distribution, conditional on the available predictor information. The generative structure of our CGM approach becomes apparent in the second part of the model, where n_{out} random samples of D_{latent} -dimensional latent variables are drawn independently from a standard multivariate normal distribution, that is,

$$\mathbf{z}_1, \dots, \mathbf{z}_{n_{\text{out}}} \sim \mathcal{N}^{D_{\text{latent}}}(\mathbf{0}, \mathbf{I}),$$

where $\mathbf{z}_i \in \mathbb{R}^{D_{\text{latent}}}$ for $i = 1, \dots, n_{\text{out}}$. The number of samples n_{out} generated in this step directly controls the final number of output samples obtained from the multivariate CGM forecast distribution and thus allows for generating an arbitrary number of postprocessed samples.⁷ In a next step, the noise encoder module aims to encode the uncertainty from the ensemble weather forecasts into the latent distribution by adjusting the scale of the latent variables via a linear mapping. To that end, we first model the conditional variance of the latent distribution via a fully connected NN

$$\boldsymbol{\delta} = h^{\delta}(\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(D)}),$$

where an exponential activation function is applied in the output layer to ensure positivity of the variances $\boldsymbol{\delta} \in \mathbb{R}^{D_{\text{latent}}}$. The scaling coefficient vector $\boldsymbol{\delta} = [\delta_1, \dots, \delta_{D_{\text{latent}}}]^T$ is then used to adjust the variance of the latent noise variables, that is,

$$\tilde{\mathbf{z}}_i = \boldsymbol{\delta} \odot \mathbf{z}_i$$

for $i = 1, \dots, n_{\text{out}}$.⁸

In the final part of the model, the scale-adjusted latent variables from the noise encoder module are now combined with additional predictors to yield the final conditional noise component $\mathbf{y}_i^{\text{noise}}$. This noise decoder module is a fully connected NN

$$\mathbf{y}_i^{\text{noise}} = h^{\text{noise}}(\bar{\mathbf{X}}^{(1)}, \dots, \bar{\mathbf{X}}^{(D)}, \mathbf{s}^{(1)}, \dots, \mathbf{s}^{(D)}, \mathbf{loc}^{(1)}, \dots, \mathbf{loc}^{(D)}, \text{doy}, \tilde{\mathbf{z}}_i).$$

The final output of the model is given by the collection of realizations from the multivariate forecast distribution

$$\hat{\mathbf{y}}_i = \mathbf{y}^{\text{mean}} + \mathbf{y}_i^{\text{noise}},$$

for $i = 1, \dots, n_{\text{out}}$ for temperature. For wind speed, we additionally apply a softplus activation function

$$(5) \quad \hat{\mathbf{y}}_i = \log(1 + \exp(\mathbf{y}^{\text{mean}} + \mathbf{y}_i^{\text{noise}}))$$

to ensure nonnegativity of the obtained wind speed forecasts.

⁷Note that the description in the following will focus on a single sample \mathbf{z}_i . During the model training process, described in more detail below, we generate n_{out} independent CGM predictions.

⁸Note that we now have $\tilde{\mathbf{z}}_i \sim \mathcal{N}(\mathbf{0}, \text{diag}(\boldsymbol{\delta}))$, that is, it is an isotropic Gaussian with a variance conditional on the ensemble standard deviations. Scale adjustment approaches that have been applied in the verification of gridded forecasts in the meteorological literature (Bouallegue et al. (2020)) can be viewed as conceptually related.

Algorithm 1: CGM training algorithm

Input : data $\{(X, y)_n\}_{n=1}^N$, initial model parameters θ_0 , number of samples n_{train} , number of batches B , learning rate η

Output: Model parameters θ^*

- 1 **for** N_{epochs} **do**
- 2 Get mini batch $\{(X, y)_b\}_{b=1}^B$;
- 3 For each sample b generate a set of n_{train} random noise samples $\{z_b^1, \dots, z_b^{n_{\text{train}}}\}_{b=1}^B$
- 4 **for** $b = 1, \dots, B$ **do**
- 5 **for** $s = 1, \dots, n_{\text{train}}$ **do**
- 6 Compute forward pass $\hat{y}_b^s \leftarrow g_{\theta}(X_b, z_b^s)$
- 7 **end**
- 8 **end**
- 9 Compute loss over batch $L \leftarrow \frac{1}{B} \sum_b \text{ES}(y_b, [\hat{y}_b^1, \dots, \hat{y}_b^{n_{\text{train}}}]$);
- 10 Compute gradient $\nabla_{\theta} L$;
- 11 Update learning rate η using ADAM;
- 12 Update model parameters $\theta \leftarrow \theta - \eta \nabla_{\theta} L$
- 13 **end**

During training and for each training example (X_n, y_n) , where X_n and y_n represent the inputs and true observations at forecast case $n = 1, \dots, N$, respectively, we generate n_{out} independent predictions $\hat{y}_1, \dots, \hat{y}_{n_{\text{out}}}$ by querying the model n_{out} times. Each time, the model generates predictions from a different sampled noise vector $z_i, i = 1, \dots, n_{\text{out}}$, but uses the same inputs X_n . The training procedure is formalized in Algorithm 1.⁹

The CGM parameters, that is, the weights and biases of $h^{\text{mean}}, h^{\delta}$ and h^{noise} , are estimated by optimizing the energy score (see Section 5.1) as a loss function tailored to the specific situation of multivariate probabilistic forecasting based on an implicit representation of the forecast distribution in the form of a sample of size $n_{\text{train}} = 50$. We follow common practice in the machine learning literature and generate an ensemble of CGMs by repeating the estimation process multiple times from different random initializations to account for the randomness of the training process based on stochastic gradient descent methods (Lakshminarayanan, Pritzel and Blundell (2017), Schulz and Lerch (2022b)). Unless indicated otherwise, we will generate a set of $n_{\text{out}} = 50$ multivariate samples to ensure comparability with the benchmark methods when making predictions on the test set and do so by repeating the model estimation 10 times and generating five samples each.

4.4. Implementation details and hyperparameter choices. Multiple hyperparameters need to be determined for the CGM implementation. For the specific setup of the individual modules of the model ($h^{\text{mean}}, h^{\delta}$ and h^{noise}), we use two hidden layers with 100 nodes in the noise decoder module h^{noise} for both target variables. As for the noise encoder module h^{δ} , we use two hidden layers with 100 nodes each for wind speed but one linear dense layer for temperature. Initial experiments indicated no substantial changes in the resulting model performance, we thus do not further optimize this component of the model architecture. Note that we also tested a nonlinear model, based on a fully connected NN for y^{mean} , which did not lead to substantial improvements. Details are provided in the Supplementary Material. The

⁹Note that this procedure might be slow on a CPU, but the innermost loop in Algorithm 1 is fully parallelizable and can thus be efficiently implemented on GPUs.

exponential linear unit (ELU; Clevert, Unterthiner and Hochreiter (2015)) activation function is applied in all hidden layers.

The number of latent variables, D^{latent} , the learning rate and the batch size are initially determined using the distributed asynchronous hyperparameter optimization technique from Bergstra, Yamins and Cox (2013) implemented in the `hyperopt` package. Based on the automated hyperparameter optimization and selected additional experiments, we set D^{latent} to 10, use a learning rate of 0.001 and a batch size of 64 for both target variables. The results are relatively robust to changes in these hyperparameters. Exemplary results are illustrated in the form of ablation studies in the Supplementary Material. The model is trained using stochastic gradient descent optimization, based on the Adam optimizer (Kingma and Ba (2014)), where we employ an early stopping criterion with a patience of 10 epochs to avoid overfitting. The maximum number of epochs is set to 300, but generally, the CGM training and validation losses converge after a few epochs.

All available predictors (listed in Table 1), except for the target variable and date information, are normalized by removing the mean and scaling to unit variance using the standard scaler from the `scikit-learn` package (Pedregosa et al. (2011)).

5. Results. We here compare the CGM predictions with the various benchmark methods, based on two-step procedures of separately modeling marginal distributions and multivariate dependencies, introduced in Section 3.2, that is, EMOS+ECC, EMOS+GCA, DRN+ECC and DRN+GCA. To do so, we first provide some background information on forecast evaluation methods (Section 5.1) and present the general setup of our experiments (Section 5.2) and univariate results (Section 5.3). The main focus is on the multivariate forecast performance presented in Sections 5.4–5.6.

5.1. Forecast evaluation methods. We briefly review key concepts relevant to the evaluation of probabilistic forecasts and refer to Rasp and Lerch ((2018), Appendix A) and Lerch et al. ((2020), Appendix B) for details. The general goal of probabilistic forecasting is to maximize the sharpness of a predictive distribution subject to calibration (Gneiting, Balabdaoui and Raftery (2007)). In order to assess calibration and sharpness simultaneously, proper scoring rules are now widely used for comparative evaluation of probabilistic forecasts. A scoring rule $S(F, y)$ assigns a numerical score to a pair of a predictive distribution $F \in \mathcal{F}$ and a realizing observation $y \in \Omega$, where \mathcal{F} is a class of probability distributions on Ω . It is called proper if the true distribution of the observation minimizes the expected score, that is, if $\mathbb{E}S(G, Y) \leq \mathbb{E}S(F, Y)$ if $Y \sim G$ for all $F, G \in \mathcal{F}$ (Gneiting and Raftery (2007a)). The continuous ranked probability score (CRPS; Matheson and Winkler (1976)), given by

$$\text{CRPS}(F, y) = \int_{-\infty}^{\infty} (F(z) - \mathbb{1}\{z \geq y\})^2 dz,$$

where $\mathbb{1}$ denotes the indicator function and F is assumed to have a finite first moment, is a popular proper scoring rule for univariate probabilistic forecasts (i.e., $\Omega \subset \mathbb{R}$). Closed-form analytical expressions are available for many parametric forecast distributions and probabilistic forecasts given in the form of a simulated sample (Jordan, Krüger and Lerch (2019)).

While the definition of proper scoring rules can, in principle, be straightforwardly extended toward multivariate settings with $\Omega \subset \mathbb{R}^D$, many practical questions remain open, and a variety of multivariate proper scoring rules have been proposed over the past years (Petropoulos et al. (2022)). Most of these multivariate proper scoring rules focus on multivariate probabilistic forecasts in the form of samples from the forecast distributions.

Using notation introduced in Section 3, the energy score (ES; Gneiting and Raftery (2007a)),

$$\text{ES}(F, \mathbf{y}) = \frac{1}{M} \sum_{i=1}^M \|\mathbf{X}_i - \mathbf{y}\| - \frac{1}{2M^2} \sum_{i=1}^M \sum_{j=1}^M \|\mathbf{X}_i - \mathbf{X}_j\|,$$

where $\|\cdot\|$ is the Euclidean norm on \mathbb{R}^D , and the variogram score of order p (VS^p ; Scheuerer and Hamill (2015)),

$$\text{VS}^p(F, \mathbf{y}) = \sum_{i=1}^D \sum_{j=1}^D w_{i,j} \left(|y^{(i)} - y^{(j)}|^p - \frac{1}{M} \sum_{k=1}^M |X_k^{(i)} - X_k^{(j)}|^p \right)^2,$$

are the most popular examples of multivariate proper scoring rules. In the definition of the VS, $w_{i,j} \geq 0$ is a nonnegative weight for pairs of component combinations, and p is the order of the VS. We use an unweighted (i.e., $w_{i,j} = 1$ for all i, j) version of the VS with order $p = 0.5$ throughout, following suggestions of Scheuerer and Hamill (2015) and utilizing implementations provided in Jordan, Krüger and Lerch (2019). Other multivariate proper scoring rules have been proposed, including copula scores that focus on the dependence structure (Ziel and Berk (2019)), and weighted versions of ES and VS (Allen, Ginsbourger and Ziegel (2022)). We present additional results for some of these scores in the Supplementary Material.

To compare forecasting methods based on a proper scoring rule with respect to a benchmark, we will often calculate the associated skill score. With the mean score of the forecasting method of interest over a test dataset, \bar{S}_f , the corresponding mean scores of the benchmark, \bar{S}_{ref} , and the (hypothetical) optimal forecast, \bar{S}_{opt} , the skill score SS_f is calculated via

$$SS_f = \frac{\bar{S}_{\text{ref}} - \bar{S}_f}{\bar{S}_{\text{ref}} - \bar{S}_{\text{opt}}}.$$

For the scoring rules considered below, $S_{\text{opt}} = 0$. Skill scores are positively oriented with a maximum value of 1, values of 0 indicating no improvement over the benchmark and negative values indicating a worse predictive performance than the benchmark.

5.2. Setup of the multivariate postprocessing experiments. To evaluate the multivariate forecast performance of different postprocessing methods, we repeatedly subsample the station datasets described in Section 2. Focusing on spatial dependencies over geographically close stations in a setting that aims to mimic practical applications, we fix a number of dimensions $D \in \{5, 10, 20\}$ and proceed as follows.

We randomly pick a station and then select the $(D - 1)$ stations that are geographically closest to obtain a set of D stations, based on which we implement the multivariate postprocessing methods, as described in Sections 3 and 4. Next, we apply the scoring rules, introduced above, to obtain corresponding mean scores over the test set, that is, data from the calendar year 2016, and compute the corresponding skill scores.

To account for uncertainties, the above procedure is repeated 100 times for both temperature and wind speed. For skill score computations, EMOS+ECC is used as a reference method throughout. For all methods, we generate 50 multivariate samples from all postprocessed forecast distributions to ensure consistency and allow for a fair comparison.

5.3. Univariate results. While the focus of our study is on multivariate postprocessing, the univariate predictive performance constitutes an important component of the overall forecast quality. Therefore, we first focus on the univariate, marginal predictions of different postprocessing methods to investigate the performance of our CGM approach in this setting.




FIG. 3. Boxplots of mean CRPS values of different multivariate postprocessing methods with $D = 5$, including the scores of raw ensemble forecasts. The scores are based on 242 unique stations in case of temperature and 178 unique stations in case of wind speed.

To evaluate the univariate forecast performance in the experimental setup described above, we restrict our attention to the experiments for $D = 5$ and compute the mean CRPS over unique sets of stations present in the 100 repetitions of the subsampled station datasets. In case a station occurs multiple times in the randomly selected sets of stations, we only use data from the first occurrence within the 100 repetitions. We compare the univariate performance of our CGM to the raw ensemble forecasts and two univariate postprocessing approaches discussed earlier, that is, EMOS and DRN. Figure 3 shows boxplots of the mean CRPS values over the corresponding unique sets of stations for temperature and wind speed. As expected, all univariate postprocessing methods notably improve the forecast performance over the raw ensemble predictions, and the variability among different stations is reduced. For temperature the CGM forecasts generally outperform the EMOS predictions and are slightly worse than the DRN postprocessed forecasts, where the mean CRPS over all stations is improved from 0.89 for EMOS to 0.76 for DRN and 0.79 for CGM. For wind speed the CGM provides the overall best forecasts and clearly outperforms the DRN approach, with a mean CRPS improved from 0.62 for EMOS and 0.58 for DRN to 0.54 for CGM. The differences in performance can potentially be explained by the better fit of the chosen Gaussian parametric distribution for temperature, which favors DRN, whereas the choice is less clear for wind speed, favoring the nonparametric CGM approach. Additional results are provided in the Supplementary Material, including an assessment of calibration, which indicates that all postprocessing methods generally provide relatively well-calibrated forecasts.

5.4. *Multivariate results.* We now turn to the key part of our results and compare the multivariate performance of our CGM approach to the two-step postprocessing approaches used as benchmark methods. Table 2 summarizes the mean scores of different multivariate postprocessing models for temperature and wind speed. For both target variables, all postprocessing methods clearly improve the raw ensemble predictions. In comparison to the EMOS-based models, the DRN-based models show clear improvements in the multivariate performance. Regarding the choice of the reordering method, ECC and GCA lead to similar results in terms of the ES, but the GCA-based forecasts lead to better performance in terms of the VS. The

TABLE 2

Mean multivariate scores of different multivariate postprocessing methods for temperature and wind speed, averaged over the 100 repetitions of the simulation experiment

Variable	Score	D	Raw ens.	EMOS+ ECC	EMOS+ GCA	DRN+ ECC	DRN+ GCA	CGM
Temperature	ES	5	2.81	2.27	2.27	1.97	1.97	1.97
		10	4.22	3.37	3.37	2.91	2.90	2.91
		20	6.09	4.87	4.87	4.21	4.22	4.26
	VS	5	8.22	4.81	4.36	4.12	3.74	3.50
		10	39.0	22.6	21.0	19.5	18.0	16.9
		20	153	96.7	92.8	85.0	80.7	77.8
Wind speed	ES	5	2.44	1.69	1.68	1.56	1.55	1.44
		10	3.67	2.55	2.53	2.31	2.30	2.16
		20	5.04	3.52	3.51	3.23	3.22	3.04
	VS	5	9.49	4.37	4.00	4.01	3.66	3.31
		10	39.7	20.2	19.0	18.0	16.9	15.4
		20	153	82.5	78.9	75.6	72.3	67.0

CGM consistently provides the best multivariate forecasts and outperforms the state-of-the-art approaches across the variables, dimensions and evaluation metrics. The only exception to this observation are temperature forecasts evaluated with the ES, where the DRN+ECC and DRN+GCA models provide slightly better forecasts for $D = 10$ and $D = 20$. The values of all considered multivariate scoring rules increase with the spatial dimension D , which is to be expected from the definition of the scoring rules and consistent with findings in the extant literature.

To investigate the variability across the selected sets of stations, Figures 4 and 5 show boxplots of the multivariate skill scores for temperature and wind speed, respectively, using EMOS+ECC as reference method. For the temperature forecasts (Figure 4), the DRN-based two-step methods and the CGM provide consistent and comparable improvements over the reference in terms of the ES. The relative improvements in terms of the VS show a larger variability across the sets of stations and indicate a superior performance of the CGM forecasts. Among the considered two-step approaches, applying GCA leads to improvements over ECC in terms of the VS but similar results in terms of the ES. The above observations apply to all considered spatial dimensions, and we do not observe any obvious trends in terms of D , indicating that consistent improvements can be observed also in the higher-dimensional settings in the experiments.

Qualitatively, similar results can be observed for the multivariate wind speed forecasts shown in Figure 5. The main difference to the results for temperature are the notably larger improvements of the CGM forecasts in comparison to the DRN-based approaches, particularly in terms of the ES. Interestingly, in terms of the VS at $D = 5$, the DRN+ECC models here fail to outperform the EMOS+GCA forecasts despite the incorporation of additional predictor variables in the marginal distributions. A potential explanation for this observation is that the disadvantages, due to the misspecifications in the multivariate dependence structure of the raw ensemble forecasts, which serve as a dependence template for ECC, outweigh the benefits of incorporating additional predictors in the marginal distributions. Similar to the temperature forecasts, the DRN+GCA model results in better forecasts than the DRN+ECC approach but performs notably worse than the CGM. Additional verification results, including assessments of multivariate calibration, significance tests on score differences, and results for other multivariate proper scoring rules, are available in the Supplementary Material.

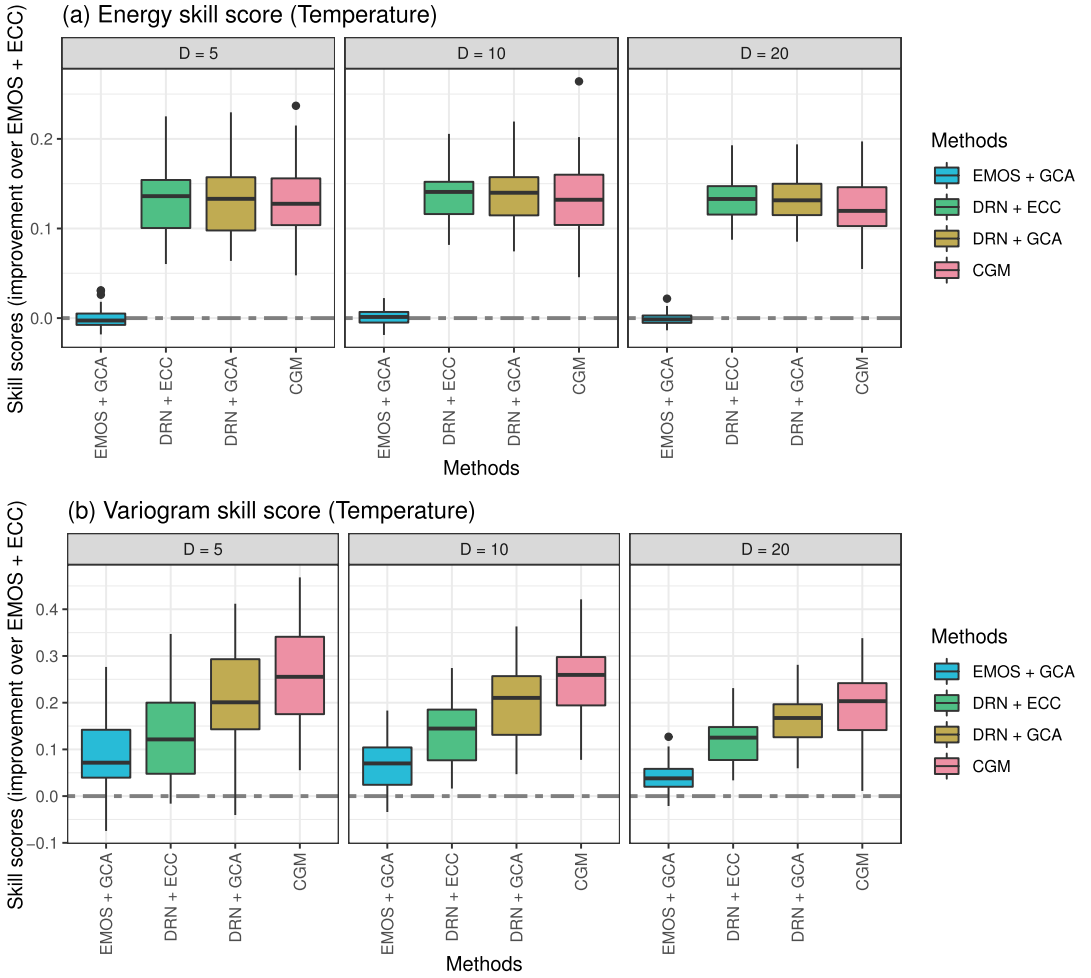


FIG. 4. Boxplots of: (a) energy skill scores and (b) variogram skill scores of different multivariate postprocessing methods for temperature across the 100 repetitions of the experiment with different sets of stations. EMOS+ECC is used as reference forecast in both cases.

5.5. *The role of additional inputs for the CGM predictive performance.* To assess the importance of incorporating additional input features for the CGM performance, Figure 6 includes a CGM variant which only uses ensemble forecasts of the target variable as input. While this CGM variant has access to the same information as the EMOS+ECC and EMOS+GCA models, it generally shows superior predictive performance, indicating improvements of the CGM approach beyond utilizing additional input features only. While the CGM variant without additional inputs typically fails to achieve forecast performance comparable to the DRN-based models, it, on average, outperforms the DRN+ECC model for wind speed in terms of the variogram score.

5.6. *CGM sample size.* In addition to incorporating arbitrary predictor variables, a particular advantage of the proposed CGM approach over ECC is that the generative procedure allows for generating an arbitrary number of samples from the predictive distribution instead of being limited to the number of ensemble members. To investigate the effect of the size n_{out} of the generated CGM ensemble, Figure 7 shows boxplots of the multivariate skill scores as functions of the ensemble size, based on the 100 repetitions of the experiment for $D = 10$. As before, we repeat the model estimation procedure of the CGM approach 10 times and generate $\frac{n_{out}}{10}$ samples each time to obtain a final postprocessed ensemble of size n_{out} .

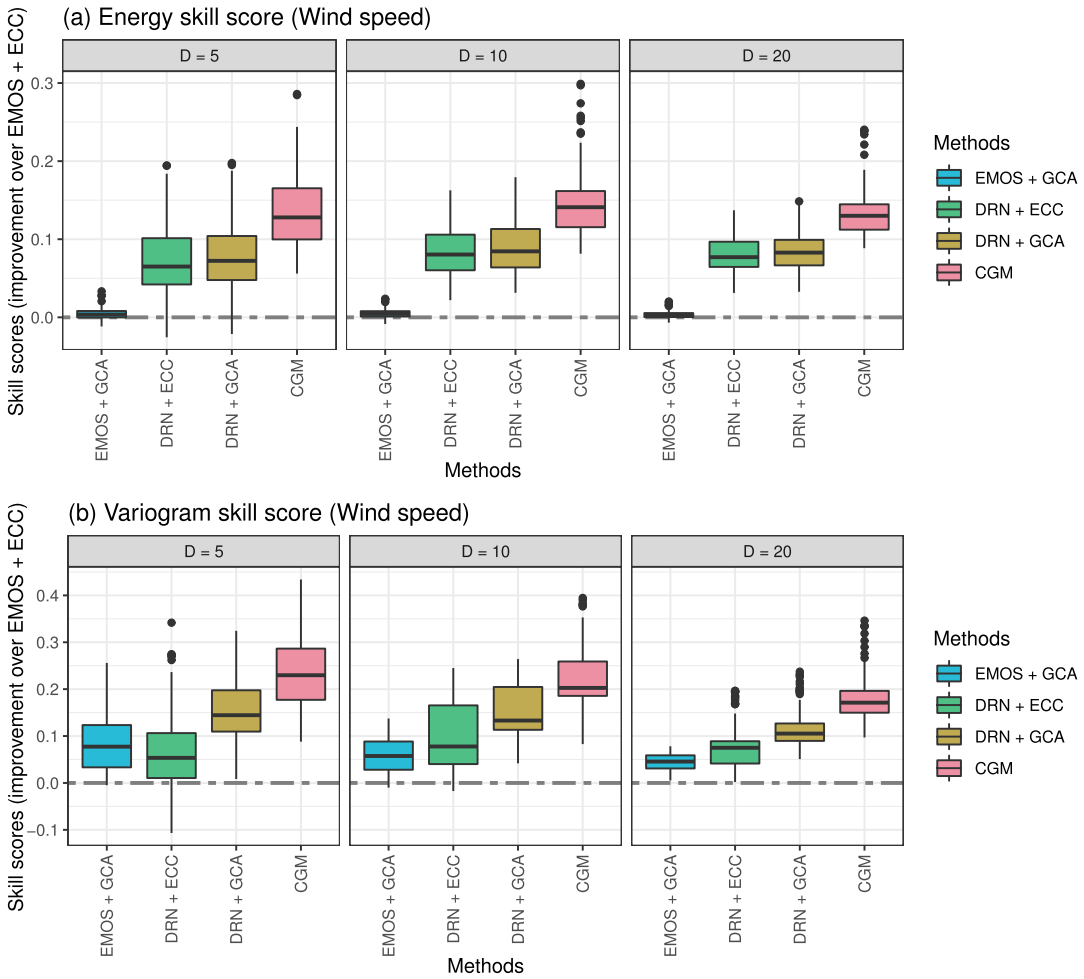


FIG. 5. As Figure 4, but for wind speed.

Compared to the reference setting of 50 CGM ensemble members, generating a larger sample from the postprocessed distributions generally improves the predictive performance, with median improvements in terms of the energy and variogram score of up to around 1.5%. The median skill score values increase notably up to an ensemble size of 200, after which some minor improvements can be observed. Additional results on other values of D are provided in the Supplementary Material.

Given a fixed CGM ensemble size n_{out} , various ensembling strategies for obtaining these n_{out} forecasts could be devised. For example, to obtain 50 CGM members, one could repeat the CGM model estimation 50 (or 25, 10, five, two, one) time(s) and generate one (or two, five, 10, 25, 50) sample(s) each, respectively. While we found that, in general, increasing the number of model runs leads to larger improvements in predictive performance, compared to increasing the number of generated samples, this needs to be balanced against the added computational costs for repeating the CGM estimation. More details on the effects of different ensembling strategies are provided in the Supplementary Material.

6. Discussion and conclusions. We propose a nonparametric multivariate postprocessing method, based on a conditional generative machine learning model, which is trained by optimizing a suitable multivariate proper scoring rule. In our CGM approach, an arbitrary number of samples from the multivariate forecast distribution is directly obtained as

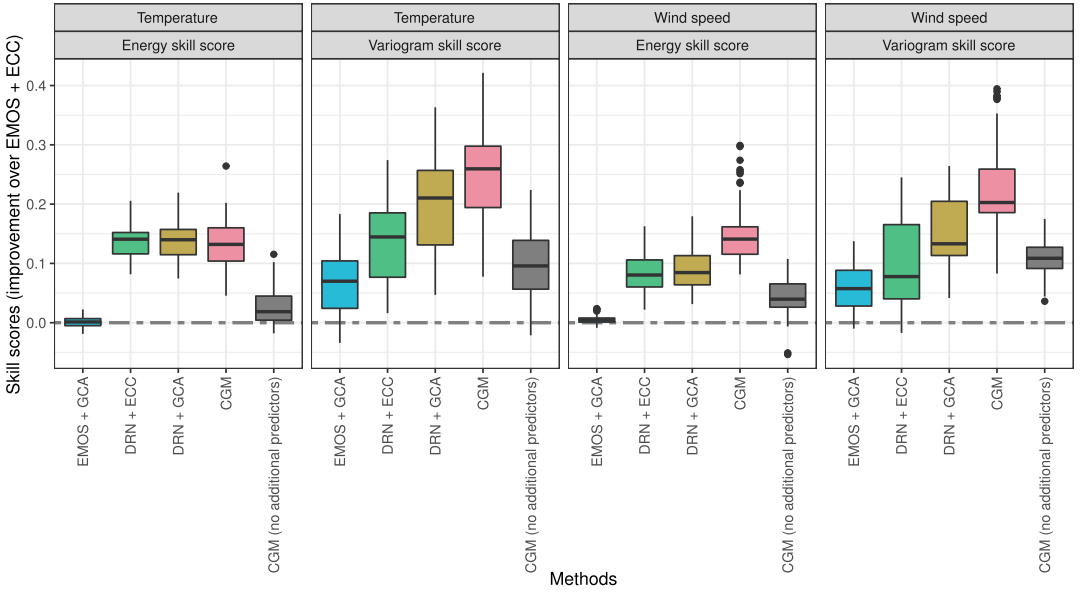


FIG. 6. Box plots of energy skill scores and variogram skill scores of different multivariate postprocessing methods, analogous to Figures 4 and 5 for $D = 10$, but including a CGM variant without additional inputs. EMOS+ECC is used as reference forecast throughout.

output of a generative deep neural network which allows for incorporating arbitrary input predictors beyond ensemble predictions of the target variables only. By circumventing the two-step structure of the state-of-the-art multivariate postprocessing approaches, the generative model aims to simultaneously correct systematic errors in the marginal distributions and the multivariate dependence structure. By contrast to the standard two-step methods, our CGM approach does not require the choice of parametric models. Furthermore, our CGM architecture can be specifically tailored to incorporate relevant exogenous information and domain knowledge in the different components of the target distribution. For example, our noise encoder module allows for dynamically reparametrizing the latent distributions of the

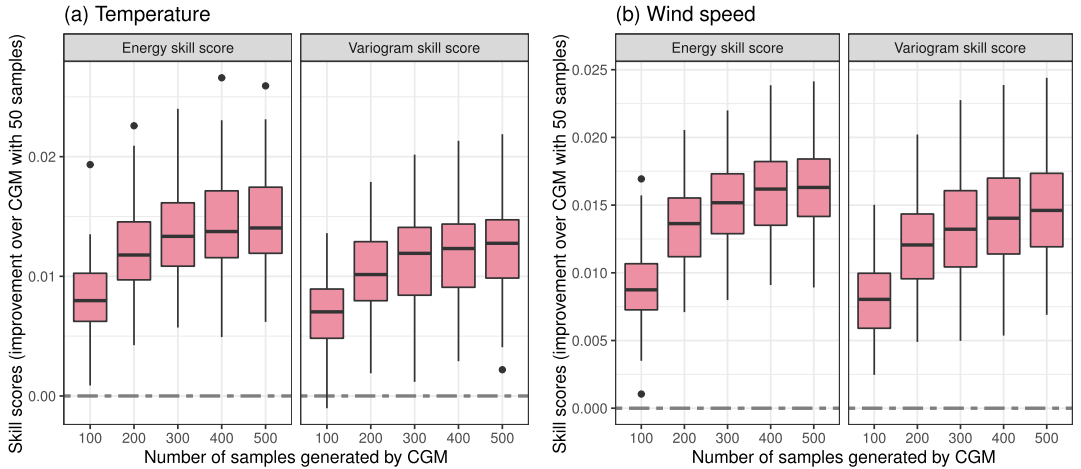


FIG. 7. Boxplots of energy skill scores and variogram skill scores of CGM forecasts with different numbers of samples generated from the multivariate postprocessed forecast distribution for (a) temperature and (b) wind speed over the 100 repetitions of the experiment with different sets of stations. The CGM approach with $n_{\text{out}} = 50$ is used as reference forecast and we only consider the case $D = 10$ here.

generative model, conditional on the standard deviations of the NWP ensemble forecasts to efficiently propagate uncertainty information.

In two case studies on spatial dependencies of temperature and wind speed forecasts at weather stations over Germany, our generative model outperforms state-of-the-art two-step methods for multivariate postprocessing where univariate postprocessing via DRN models is combined with ECC and GCA. Our CGM approach provides improvements in terms of the univariate forecast performance at individual stations and produces the best overall multivariate forecasts in terms of the energy score and the variogram score. The observed score differences are statistically significant for a large fraction of the random repetitions of the experiments, even when compared to the best-performing benchmark methods with details provided in the Supplementary Material. Overall, there are no clear differences in the performance across the considered spatial dimensions of five, 10 and 20 stations, indicating that the CGM approach works well also in higher-dimensional settings. Regarding the two target variables, we observed more pronounced improvements over the state-of-the-art two-step methods for wind speed, potentially mainly due to larger improvements in the univariate forecast performance. In terms of the two considered multivariate proper scoring rules, the relative improvements are generally larger in terms of the variogram score, indicating that our CGM approach particularly succeeds in better modeling the multivariate dependence structure. The only case where we did not observe notable differences to the performance of the benchmark methods were the results for temperature in terms of the energy score.

The clear improvements in terms of the predictive performance are likely due to key conceptual advantages of our conditional generative models over the two-step approaches, in particular their ability to incorporate arbitrary predictor variables in both the modeling of the marginal distributions and the multivariate dependencies. CGM architectures, without additional predictors beyond ensemble forecasts of the target variables, reached a better multivariate predictive performance than EMOS+ECC and EMOS+GCA models but failed to outperform the DRN-based models.

Two minor disadvantages of the CGM approach are, on the one hand, given by the slightly increased variability across the random repetitions of the experiments, due to the generative procedure, which can lead to single outliers with a worse predictive performance. Generating ensembles of CGM predictions can help to alleviate this, in particular, with an increased number of subensembles and sample size (Schulz and Lerch (2022b)). Details on different strategies for generating CGM ensembles are discussed in the Supplementary Material. On the other hand, the CGM approach is conceptually somewhat simpler than the two-step methods in that it does not require any parametric assumptions on univariate forecast distributions or multivariate dependencies and produces forecasts in a single step only; however, the computational costs of model training are larger. That said, the computational costs of CGM for multivariate postprocessing are still negligible, compared to the computational costs of generating the raw ensemble forecasts, and will not be a limiting factor in research or operations. For example, for a fixed set of $D = 20$ stations, estimating an ensemble of 10 CGMs and generating five samples each takes around two minutes on a Nvidia RTX A5000 GPU.

Over the past years, many techniques have been developed in order to better interpret and understand what machine learning methods have learned, in particular, for NNs (see McGovern et al. (2019), for an overview from a meteorological perspective). Methods from interpretable machine learning have been applied in the literature on univariate postprocessing (Rasp and Lerch (2018), Schulz and Lerch (2022a), Taillardat et al. (2016)), but the problem is more involved for multivariate postprocessing, in particular, for the generative models proposed here. While there has been some progress for GANs (Adel, Ghahramani and Weller (2018), Chen et al. (2016)), interpretation is challenging for generative models (Zhou (2022)), and the application of standard methods, such as permutation feature importance, is not straightforward.

Our results provide several avenues for further generalization and analysis. While we have focused on spatial dependencies across observation stations, it would be interesting to investigate the performance of the CGM approach on a gridded dataset. Motivated by the potential of score-based generative models to achieve comparable performance to GANs in image generation tasks (Song and Ermon (2020)), applications to multivariate postprocessing, similar to the GAN models proposed in Dai and Hemri (2021), constitute a natural starting point for future work. The CGM approach could further be applied to model temporal or intervariable dependencies. In addition, while the focus of our case studies was on the multivariate forecast performance, the results presented in Section 5.3 indicate that the univariate CGM forecasts show competitive performance, even with state-of-the-art NN-based postprocessing models applied for the marginal distributions, despite being trained in a multivariate setting. Therefore, it would also be interesting to investigate the potential of the generative models for univariate probabilistic forecasting in more detail, ideally in conjunction with considering theoretical aspects, such as the effects of choosing different proper scoring rules for optimization (Pacchiardi et al. (2021)). Furthermore, the CGM architecture could be combined with additional predictors that aim to incorporate information from flow-dependent large-scale spatial structures in the raw forecast fields. For example, Lerch and Polsterer (2022) propose convolutional autoencoder NNs to learn low-dimensional representations of spatial forecast fields, which could be used as additional CGM inputs to achieve a spatially-informed modeling of multivariate dependencies. Finally, as an alternative two-step strategy for multivariate postprocessing, it is also possible to employ generative ML methods to learn conditional copula functions (Janke, Ghanmi and Steinke (2021)) in the second step, which allow for incorporating arbitrary additional predictors.

The evaluation of multivariate probabilistic forecasts continues to represent an important methodological challenge, despite relevant recent work on multivariate proper scoring rules (Alexander et al. (2022), Allen, Ginsbourger and Ziegel (2022), Ziel and Berk (2019)). Regarding the evaluation of the CGM forecasts, the question on how to best differentiate between the contributions of improvements in univariate and multivariate components of the overall forecast performance measured via multivariate proper scoring rules is a particular challenge. Another important aspect is the evaluation of multivariate extreme events (Lerch et al. (2017)), where recent work from Allen, Ginsbourger and Ziegel (2022) could serve as a starting point for systematically investigating the effect of the sample size of our CGM and alternative approaches on the ability of the postprocessing models to provide reliable and accurate multivariate predictions of extreme events.

Acknowledgments. We thank Nina Horat, Benedikt Schulz and Tilmann Gneiting for helpful discussions, Sam Allen for providing code for the weighted multivariate scoring rules and three anonymous reviewers for insightful and constructive comments.

Funding. The research leading to these results has been done within the Young Investigator Group “Artificial Intelligence for Probabilistic Weather Forecasting” and funded by the Vector Stiftung. In addition, this project has received funding from the KIT Center for Mathematics in Sciences, Engineering and Economics under the seed funding programme.

SUPPLEMENTARY MATERIAL

Supplement to “Generative machine learning methods for multivariate ensemble postprocessing” (DOI: [10.1214/23-AOAS1784SUPPA](https://doi.org/10.1214/23-AOAS1784SUPPA); .pdf). Ablation studies regarding the architecture and hyperparameter choices of the conditional generative model and some results not shown in the paper are provided.

Code and data (DOI: [10.1214/23-AOAS1784SUPPB](https://doi.org/10.1214/23-AOAS1784SUPPB); .zip). Code and data to reproduce the postprocessed forecasts, with instructions in README.md.

REFERENCES

- ADEL, T., GHAHRAMANI, Z. and WELLER, A. (2018). Discovering interpretable representations for both deep generative and discriminative models. In *Proceedings of the 35th International Conference on Machine Learning* **80** 50–59. PMLR, Atlanta, GA, USA.
- ALEXANDER, C., COULON, M., HAN, Y. and MENG, X. (2022). Evaluating the discrimination ability of proper multi-variate scoring rules. *Ann. Oper. Res.* 1–27.
- ALLEN, S., GINSBOURGER, D. and ZIEGEL, J. (2022). Evaluating forecasts for high-impact events using transformed kernel scores. Preprint. Available at [arXiv:2202.12732](https://arxiv.org/abs/2202.12732).
- BARAN, S. and LERCH, S. (2015). Log-normal distribution based ensemble model output statistics models for probabilistic wind-speed forecasting. *Q. J. R. Meteorol. Soc.* **141** 2289–2299.
- BARAN, S. and LERCH, S. (2016). Mixture EMOS model for calibrating ensemble forecasts of wind speed. *Environmetrics* **27** 116–130. MR3481324 <https://doi.org/10.1002/env.2380>
- BARAN, S. and MÖLLER, A. (2015). Joint probabilistic forecasting of wind speed and temperature using Bayesian model averaging. *Environmetrics* **26** 120–132. MR3324906 <https://doi.org/10.1002/env.2316>
- BARAN, S., SZOKOL, P. and SZABÓ, M. (2021). Truncated generalized extreme value distribution-based ensemble model output statistics model for calibration of wind speed ensemble forecasts. *Environmetrics* **32** e2678. MR4309587 <https://doi.org/10.1002/env.2678>
- BAUER, P., THORPE, A. and BRUNET, G. (2015). The quiet revolution of numerical weather prediction. *Nature* **525** 47–55. <https://doi.org/10.1038/nature14956>
- BERGSTRA, J., YAMINS, D. and COX, D. (2013). Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *Proceedings of the 30th International Conference on Machine Learning* **28** 115–123. PMLR, Atlanta, GA, USA.
- BOUALLEGUE, Z. B., HAIDEN, T., WEBER, N. J., HAMILL, T. M. and RICHARDSON, D. S. (2020). Accounting for representativeness in the verification of ensemble precipitation forecasts. *Mon. Weather Rev.* **148** 2049–2062.
- BOUALLEGUE, Z. B., HEPELMANN, T., THEIS, S. E. and PINSON, P. (2016). Generation of scenarios from calibrated ensemble forecasts with a dual-ensemble copula-coupling approach. *Mon. Weather Rev.* **144** 4737–4750.
- BOUGEAULT, P., TOTH, Z. et al. (2010). The THORPEX interactive grand global ensemble. *Bull. Am. Meteorol. Soc.* **91** 1059–1072.
- BREMNES, J. B. (2020). Ensemble postprocessing using quantile function regression based on neural networks and Bernstein polynomials. *Mon. Weather Rev.* **148** 403–414.
- CHALOULOS, G. and LYGEROS, J. (2007). Effect of wind correlation on aircraft conflict probability. *J. Guid. Control Dyn.* **30** 1742–1752.
- CHAPMAN, W. E., MONACHE, L. D., ALESSANDRINI, S., SUBRAMANIAN, A. C., RALPH, F. M., XIE, S.-P., LERCH, S. and HAYATBINI, N. (2022). Probabilistic predictions from deterministic atmospheric river forecasts with deep learning. *Mon. Weather Rev.* **150** 215–234.
- CHEN, J., JANKE, T., STEINKE, F. and LERCH, S. (2024). Supplement to “Generative machine learning methods for multivariate ensemble postprocessing.” <https://doi.org/10.1214/23-AOAS1784SUPPA>, <https://doi.org/10.1214/23-AOAS1784SUPPB>
- CHEN, X., DUAN, Y., HOUTHOOFT, R., SCHULMAN, J., SUTSKEVER, I. and ABBEEL, P. (2016). InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems* **29**. Curran Associates, Red Hook.
- CLARK, M., GANGOPADHYAY, S., HAY, L., RAJAGOPALAN, B. and WILBY, R. (2004). The Schaake shuffle: A method for reconstructing space–time variability in forecasted precipitation and temperature fields. *J. Hydrometeorol.* **5** 243–262.
- CLEVERT, D.-A., UNTERTHINER, T. and HOCHREITER, S. (2015). Fast and accurate deep network learning by exponential linear units (ELUs). Preprint. Available at [arXiv:1511.07289](https://arxiv.org/abs/1511.07289).
- DAI, Y. and HEMRI, S. (2021). Spatially coherent postprocessing of cloud cover ensemble forecasts. *Mon. Weather Rev.* **149** 3923–3937.
- DZIUGAITE, G. K., ROY, D. M. and GHAHRAMANI, Z. (2015). Training generative neural networks via Maximum Mean Discrepancy optimization. Preprint. Available at [arXiv:1505.03906](https://arxiv.org/abs/1505.03906).
- FELDMANN, K., SCHEUERER, M. and THORARINSDOTTIR, T. L. (2015). Spatial postprocessing of ensemble forecasts for temperature using nonhomogeneous Gaussian regression. *Mon. Weather Rev.* **143** 955–971.
- GNEITING, T., BALABDAOUI, F. and RAFTERY, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **69** 243–268. MR2325275 <https://doi.org/10.1111/j.1467-9868.2007.00587.x>
- GNEITING, T. and RAFTERY, A. E. (2007a). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* **102** 359–378. MR2345548 <https://doi.org/10.1198/016214506000001437>

- GNEITING, T., RAFTERY, A. E., WESTVELD, A. H. and GOLDMAN, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Weather Rev.* **133** 1098–1118.
- GOODFELLOW, I. J., POUGET-ABADIE, J., MIRZA, M., XU, B., WARDE-FARLEY, D., OZAIR, S., COURVILLE, A. and BENGIO, Y. (2014). Generative adversarial nets. In *NIPS'14: Proceedings of the 27th International Conference on Neural Information Processing Systems—Volume 2* 2672–2680. MIT Press, Cambridge, MA, USA.
- GRETTON, A., BORGWARDT, K. M., RASCH, M. J., SCHÖLKOPF, B. and SMOLA, A. (2012). A kernel two-sample test. *J. Mach. Learn. Res.* **13** 723–773. [MR2913716](#)
- GUI, J., SUN, Z., WEN, Y., TAO, D. and YE, J. (2021). A review on generative adversarial networks: Algorithms, theory, and applications. *IEEE Trans. Knowl. Data Eng.* 1–1.
- HARRIS, L., MCRAE, A. T. T., CHANTRY, M., DUEBEN, P. D. and PALMER, T. N. (2022). A generative deep learning approach to stochastic downscaling of precipitation forecasts. *J. Adv. Model. Earth Syst.* **14** e2022MS003120.
- HAUPT, S. E., CHAPMAN, W., ADAMS, S. V., KIRKWOOD, C., HOSKING, J. S., ROBINSON, N. H., LERCH, S. and SUBRAMANIAN, A. C. (2021). Towards implementing artificial intelligence post-processing in weather and climate: Proposed actions from the Oxford 2019 workshop. *Philos. Trans. R. Soc. A* **379** 20200091. [MR4236151](#) <https://doi.org/10.1098/rsta.2020.0091>
- HESS, P., DRÜKE, M., PETRI, S., STRNAD, F. and BOERS, N. (2022). Physically constrained generative adversarial networks for improving precipitation fields from Earth system models. *Nature Machine Intelligence* **4** 828–839. <https://doi.org/10.1038/s42256-022-00540-1>
- HU, Y., SCHMEITS, M. J., VAN ANDEL, S. J., VERKADE, J. S., XU, M., SOLOMATINE, D. P. and LIANG, Z. (2016). A stratified sampling approach for improved sampling from a calibrated ensemble forecast distribution. *J. Hydrometeorol.* **17** 2405–2417.
- JANKE, T., GHANMI, M. and STEINKE, F. (2021). Implicit generative copulas. In *Advances in Neural Information Processing Systems* **34** 26028–26039. Curran Associates, Red Hook.
- JANKE, T. and STEINKE, F. (2020). Probabilistic multivariate electricity price forecasting using implicit generative ensemble post-processing. In *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)* 1–6. IEEE, Los Alamitos.
- JORDAN, A., KRÜGER, F. and LERCH, S. (2019). Evaluating probabilistic forecasts with scoringRules. *J. Stat. Softw.* **90** 1–37.
- KINGMA, D. P. and BA, J. (2014). Adam: A method for stochastic optimization. Preprint. Available at [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- LAKATOS, M., LERCH, S., HEMRI, S. and BARAN, S. (2023). Comparison of multivariate post-processing methods using global ECMWF ensemble forecasts. *Q. J. R. Meteorol. Soc.* **149** 856–877.
- LAKSHMINARAYANAN, B., PRITZEL, A. and BLUNDELL, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems* 6405–6416. Curran Associates Inc., Red Hook, NY, USA.
- LANG, M. N., LERCH, S., MAYR, G. J., SIMON, T., STAUFFER, R. and ZEILEIS, A. (2020). Remember the past: A comparison of time-adaptive training schemes for non-homogeneous regression. *Nonlinear Process. Geophys.* **27** 23–34.
- LANG, M. N., MAYR, G. J., STAUFFER, R. and ZEILEIS, A. (2019). Bivariate Gaussian models for wind vectors in a distributional regression framework. *Adv. Stat. Climatol. Meteorol. Oceanogr.* **5** 115–132.
- LEINONEN, J., NERINI, D. and BERNE, A. (2021). Stochastic super-resolution for downscaling time-evolving atmospheric fields with a generative adversarial network. *IEEE Trans. Geosci. Remote Sens.* **59** 7211–7223.
- LERCH, S. and BARAN, S. (2017). Similarity-based semilocal estimation of post-processing models. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **66** 29–51. [MR3611676](#) <https://doi.org/10.1111/rssc.12153>
- LERCH, S., BARAN, S., MÖLLER, A., GROSS, J., SCHEFZIK, R., HEMRI, S. and GRAETER, M. (2020). Simulation-based comparison of multivariate ensemble post-processing methods. *Nonlinear Process. Geophys.* **27** 349–371.
- LERCH, S. and POLSTERER, K. L. (2022). Convolutional autoencoders for spatially-informed ensemble post-processing. International Conference on Learning Representations (ICLR) 2022—AI for Earth and Space Science Workshop. <https://doi.org/10.48550/arXiv.2204.05102>
- LERCH, S. and THORARINSDOTTIR, T. L. (2013). Comparison of non-homogeneous regression models for probabilistic wind speed forecasting. *Tellus, Ser. A Dyn. Meteorol. Oceanogr.* **65** 21206.
- LERCH, S., THORARINSDOTTIR, T. L., RAVAZZOLO, F. and GNEITING, T. (2017). Forecaster's dilemma: Extreme events and forecast evaluation. *Statist. Sci.* **32** 106–127. [MR3634309](#) <https://doi.org/10.1214/16-STS588>
- LI, Y., SWERSKY, K. and ZEMEL, R. (2015). Generative moment matching networks. In *Proceedings of the 32nd International Conference on Machine Learning* **37** 1718–1727. PMLR, Lille, France.

- MATHESON, J. E. and WINKLER, R. L. (1976). Scoring rules for continuous probability distributions. *Manage. Sci.* **22** 1087–1096.
- MCGOVERN, A., LAGERQUIST, R., GAGNE, D. J., JERGENSEN, G. E., ELMORE, K. L., HOMEYER, C. R. and SMITH, T. (2019). Making the black box more transparent: Understanding the physical implications of machine learning. *Bull. Am. Meteorol. Soc.* **100** 2175–2199.
- MOHAMED, S. and LAKSHMINARAYANAN, B. (2016). Learning in Implicit Generative Models. Preprint. Available at [arXiv:1610.03483](https://arxiv.org/abs/1610.03483).
- MÖLLER, A., LENKOSKI, A. and THORARINSDOTTIR, T. L. (2013). Multivariate probabilistic forecasting using ensemble Bayesian model averaging and copulas. *Q. J. R. Meteorol. Soc.* **139** 982–991.
- MUSCHINSKI, T., LANG, M. N., MAYR, G. J., MESSNER, J. W., ZEILEIS, A. and SIMON, T. (2022). Predicting power ramps from joint distributions of future wind speeds. *Wind Energy Sci.* **7** 2393–2405.
- NELSEN, R. B. (2006). *An Introduction to Copulas*, 2nd ed. *Springer Series in Statistics*. Springer, New York. [MR2197664 https://doi.org/10.1007/s11229-005-3715-x](https://doi.org/10.1007/s11229-005-3715-x)
- PACCHIARDI, L., ADEWOYIN, R., DUEBEN, P. and DUTTA, R. (2021). Probabilistic forecasting with generative networks via scoring rule minimization. Preprint. Available at [arXiv:2112.08217](https://arxiv.org/abs/2112.08217).
- PANTILLON, F., LERCH, S., KNIPPERTZ, P. and CORSMEIER, U. (2018). Forecasting wind gusts in winter storms using a calibrated convection-permitting ensemble. *Q. J. R. Meteorol. Soc.* **144** 1864–1881.
- PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A. et al. (2011). Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12** 2825–2830. [MR2854348](https://doi.org/10.48550/jmlr.2011.12.2825)
- PENNINGTON, J., SOCHER, R. and MANNING, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* 1532–1543. Association for Computational Linguistics, Doha, Qatar.
- PERRONE, E., SCHICKER, I. and LANG, M. N. (2020). A case study of empirical copula methods for the statistical correction of forecasts of the ALADIN-LAEF system. *Meteorol. Z.* **29** 277–288.
- PETROPOULOS, F. et al. (2022). Forecasting: Theory and practice. *Int. J. Forecast.* **38** 705–871.
- PINSON, P. and GIRARD, R. (2012). Evaluating the quality of scenarios of short-term wind power generation. *Appl. Energy* **96** 12–20.
- PINSON, P. and MESSNER, J. W. (2018). Chapter 9—Application of postprocessing for renewable energy. In *Statistical Postprocessing of Ensemble Forecasts* 241–266. Elsevier, Amsterdam.
- PRICE, I. and RASP, S. (2022). Increasing the accuracy and resolution of precipitation forecasts using deep generative models. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics* 1511–1526. PMLR, Atlanta, GA, USA.
- RASP, S. and LERCH, S. (2018). Neural networks for postprocessing ensemble weather forecasts. *Mon. Weather Rev.* **146** 3885–3900.
- RAVURI, S., LENC, K., WILLSON, M. et al. (2021). Skilful precipitation nowcasting using deep generative models of radar. *Nature* **597** 672–677.
- SCHFZIK, R. (2016). A similarity-based implementation of the Schaake shuffle. *Mon. Weather Rev.* **144** 1909–1921.
- SCHFZIK, R. and MÖLLER, A. (2018). Chapter 4—Ensemble postprocessing methods incorporating dependence structures. In *Statistical Postprocessing of Ensemble Forecasts* 91–125. Elsevier, Amsterdam.
- SCHFZIK, R., THORARINSDOTTIR, T. L. and GNEITING, T. (2013). Uncertainty quantification in complex simulation models using ensemble copula coupling. *Statist. Sci.* **28** 616–640. [MR3161590 https://doi.org/10.1214/13-STS443](https://doi.org/10.1214/13-STS443)
- SCHUEERER, M. and HAMILL, T. M. (2015). Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Mon. Weather Rev.* **143** 1321–1334.
- SCHUEERER, M., HAMILL, T. M., WHITIN, B., HE, M. and HENKEL, A. (2017). A method for preferential selection of dates in the Schaake shuffle approach to constructing spatiotemporal forecast fields of temperature and precipitation. *Water Resour. Res.* **53** 3029–3046.
- SCHUEERER, M. and MÖLLER, D. (2015). Probabilistic wind speed forecasting on a grid based on ensemble model output statistics. *Ann. Appl. Stat.* **9** 1328–1349. [MR3418725 https://doi.org/10.1214/15-AOAS843](https://doi.org/10.1214/15-AOAS843)
- SCHUEERER, M., SWITANEK, M. B., WORSNOP, R. P. and HAMILL, T. M. (2020). Using artificial neural networks for generating probabilistic subseasonal precipitation forecasts over California. *Mon. Weather Rev.* **148** 3489–3506.
- SCHUHEN, N., THORARINSDOTTIR, T. L. and GNEITING, T. (2012). Ensemble model output statistics for wind vectors. *Mon. Weather Rev.* **140** 3204–3219.
- SCHULZ, B. and LERCH, S. (2022a). Machine learning methods for postprocessing ensemble forecasts of wind gusts: A systematic comparison. *Mon. Weather Rev.* **150** 235–257.
- SCHULZ, B. and LERCH, S. (2022b). Aggregating distribution forecasts from deep ensembles. Preprint. Available at [arXiv:2204.02291](https://arxiv.org/abs/2204.02291).

- SEJDINOVIC, D., SRIPERUMBUDUR, B., GRETTON, A. and FUKUMIZU, K. (2013). Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Ann. Statist.* **41** 2263–2291. MR3127866 <https://doi.org/10.1214/13-AOS1140>
- SKLAR, M. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Stat. Univ. Paris* **8** 229–231. MR0125600
- SLOUGHTER, M., GNEITING, T. and RAFTERY, A. E. (2010). Probabilistic wind speed forecasting using ensembles and Bayesian model averaging. *J. Amer. Statist. Assoc.* **105** 25–35. MR2757189 <https://doi.org/10.1198/jasa.2009.ap08615>
- SONG, Y. and ERMON, S. (2020). Improved techniques for training score-based generative models. In *NIPS'20: Proceedings of the 34th International Conference on Neural Information Processing Systems* **1043** 12438–12448. Curran Associates Inc., Red Hook, NY, USA.
- SZÉKELY, G. J. and RIZZO, M. L. (2013). Energy statistics: A class of statistics based on distances. *J. Stat. Plan. Inference.* **143** 1249–1272.
- TAILLARDAT, M., MESTRE, O., ZAMO, M. and NAVEAU, P. (2016). Calibrated ensemble forecasts using quantile regression forests and ensemble model output statistics. *Mon. Weather Rev.* **144** 2375–2393.
- THORARINSDOTTIR, T. L. and GNEITING, T. (2010). Probabilistic forecasts of wind speed: Ensemble model output statistics by using heteroscedastic censored regression. *J. Roy. Statist. Soc. Ser. A* **173** 371–388. MR2751882 <https://doi.org/10.1111/j.1467-985X.2009.00616.x>
- VANNITSEM, S., BREMNES, J. B., DEMAAYER, J., EVANS, G. R., FLOWERDEW, J., HEMRI, S., LERCH, S., ROBERTS, N., THEIS, S. et al. (2021). Statistical postprocessing for weather forecasts: Review, challenges, and avenues in a big data world. *Bull. Am. Meteorol. Soc.* **102** E681–E699.
- VAN SCHAEYBROECK, B. and VANNITSEM, S. (2015). Ensemble post-processing using member-by-member approaches: Theoretical aspects. *Q. J. R. Meteorol. Soc.* **141** 807–818.
- WILKS, D. S. (2015). Multivariate ensemble model output statistics using empirical copulas. *Q. J. R. Meteorol. Soc.* **141** 945–952.
- WORSNOP, R. P., SCHEUERER, M., HAMILL, T. M. and LUNDQUIST, J. K. (2018). Generating wind power scenarios for probabilistic ramp event prediction using multivariate statistical post-processing. *Wind Energy Sci.* **3** 371–393.
- ZHOU, B. (2022). Interpreting generative adversarial networks for interactive image generation. In *XxAI—Beyond Explainable AI: International Workshop, Held in Conjunction with ICML 2020, July 18, 2020, Vienna, Austria, Revised and Extended Papers* 167–175. Springer, Cham.
- ZIEL, F. and BERK, K. (2019). Multivariate forecasting evaluation: On sensitive and strictly proper scoring rules. Preprint. Available at [arXiv:1910.07325](https://arxiv.org/abs/1910.07325).