Generative ML methods for multivariate ensemble post-processing

ECMWF–ESA Workshop on Machine Learning for Earth Observation and Prediction

Sebastian Lerch, based on joint work with Jieyu Chen, Tim Janke and Florian Steinke | November 2022

Motivation

Ensemble forecasts typically show systematic biases and lack calibration.

Post-processing can successfully correct errors and has become a standard practice in research and operations.

Overview of the state of the art: Vannitsem et al. (2021)

Statistical Postprocessing for Weather Forecasts – Review, Challenges and Avenues in a Big Data World

Stéphane Vannitsem; John Bjørnar Bremnes; Jonathan Demaeyer; Gavin R. Evans; Jonathan Flowerdew; Stephan Hemri; Sebastian Lerch; Nigel Roberts; Susanne Theis; Aitor Atencia ... Show more

Bull. Amer. Meteor. Soc. 1–44.

https://doi.org/10.1175/BAMS-D-19-0308.1

Motivation

Karlsruhe Institute of Technology

Ensemble forecasts typically show systematic biases and lack calibration.

Post-processing can successfully correct errors and has become a standard practice in research and operations.

Overview of the state of the art: Vannitsem et al. (2021)

Statistical Postprocessing for Weather Forecasts – Review, Challenges and Avenues in a Big Data World

Stéphane Vannitsem; John Bjørnar Bremnes; Jonathan Demaeyer; Gavin R. Evans; Jonathan Flowerdew; Stephan Hemri; Sebastian Lerch; Nigel Roberts; Susanne Theis; Aitor Atencia ... Show more

Bull. Amer. Meteor. Soc. 1–44.

https://doi.org/10.1175/BAMS-D-19-0308.1

While various methods have been developed for univariate post-processing for single locations, lead times or weather variables, many applications require accurate models of multivariate dependencies.

Key examples include hydrological applications, air traffic management and energy forecasting.

Data

- ECMWF 2-day ahead 50-member ensemble forecasts (temperature & wind speed at 00 UTC)
- 16 auxiliary variables available as additional predictors + station information
- 2007–2015 for training, 2016 for evaluation
- focus: spatial dependencies across stations

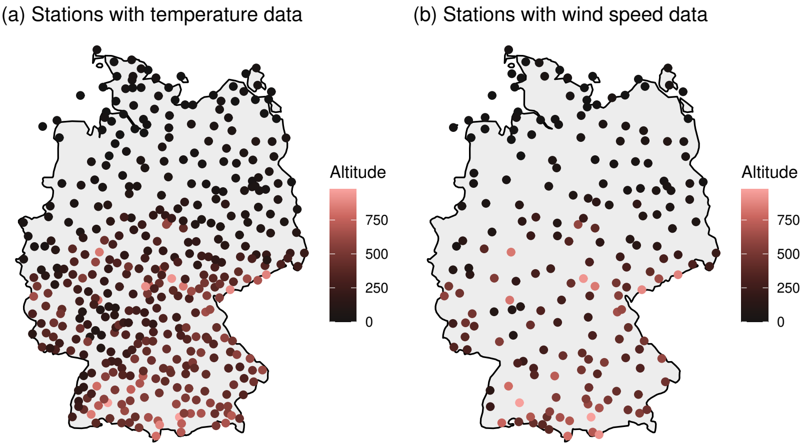

State of the art approaches for multivariate post-processing

Two-step approaches separately model univariate (marginal) distributions and multivariate dependencies:

State of the art approaches for multivariate post-processing

Two-step approaches separately model univariate (marginal) distributions and multivariate dependencies:

1. Univariate post-processing: Distributional regression models correct systematic biases and lack of calibration. A sample is generated from the post-processed distributions.
   - simple standard approach: ensemble model output statistics (EMOS, Gneiting et al., 2005, MWR)
   - state of the art: ML methods allow for incorporating additional predictors (e.g., NN-based methods, Rasp and Lerch, 2018, MWR)

Sebastian.Lerch@kit.edu

State of the art approaches for multivariate post-processing

Two-step approaches separately model univariate (marginal) distributions and multivariate dependencies:

1. Univariate post-processing: Distributional regression models correct systematic biases and lack of calibration. A sample is generated from the post-processed distributions.
   - simple standard approach: ensemble model output statistics (EMOS, Gneiting et al., 2005, MWR)
   - state of the art: ML methods allow for incorporating additional predictors (e.g., NN-based methods, Rasp and Lerch, 2018, MWR)
2. Multivariate dependencies are restored by applying copula-based methods to re-order the sample values. Prominent examples:
   - ensemble copula coupling (ECC): based on dependence structure in the raw ensemble
   - Schaake shuffle (SSh): as ECC, but based on past observations
   - Gaussian copula approach (GCA): parametric approach based on Gaussian copulas

   Comparative studies indicate that there is no consistently best approach and differences in performance are generally small.

Limitations of state of the art approaches

All state of the art methods for multivariate post-processing share common key limitations.

Most importantly, there is no straightforward way to include additional predictors beyond the target variable in the second step of restoring multivariate dependencies.

Further, the number of samples from the multivariate forecast distribution is limited by the size of the ensemble (ECC) or the historical dataset (SSh).

Limitations of state of the art approaches

All state of the art methods for multivariate post-processing share common key limitations.

Most importantly, there is no straightforward way to include additional predictors beyond the target variable in the second step of restoring multivariate dependencies.

Further, the number of samples from the multivariate forecast distribution is limited by the size of the ensemble (ECC) or the historical dataset (SSh).

We propose a novel nonparametric multivariate post-processing method based on generative ML.

Samples are directly obtained as output of a generative deep neural network which allows for

- incorporating arbitrary input predictors
- circumventing the two-step structure of traditional methods to simultaneously model distributions and multivariate dependencies
- generating an arbitrary number of samples.

Generative ML

Generative ML

Artificial, GAN-generated images of human faces,
Source: `https://thispersondoesnotexist.com/`

Generative ML

Artificial, GAN-generated images of human faces,
Source: `https://thispersondoesnotexist.com/`

Generative ML aims at learning properties underlying a training dataset to generate realistic new data.

Implicit generative models aim to provide a representation of the probability distribution of a target variable by defining a stochastic procedure to generate samples.

Generative ML

Artificial, GAN-generated images of human faces,
Source: `https://thispersondoesnotexist.com/`

Generative ML aims at learning properties underlying a training dataset to generate realistic new data.

Implicit generative models aim to provide a representation of the probability distribution of a target variable by defining a stochastic procedure to generate samples.

Impressive results have been achieved by GANs for image processing tasks, but GAN training tends to be challenging.

We aim for a conditional generative model of the form

$$\boldsymbol{Y}_i = g_\theta(\boldsymbol{X}, \boldsymbol{Z}_i), \tag{1}$$

where the discriminator part of a GAN is replaced by a suitable multivariate evaluation metric (the energy distance).

Multivariate samples $\boldsymbol{Y}_i$ are generated as the output of a NN, conditional on ensemble forecasts $\boldsymbol{X}$ and random noise $\boldsymbol{Z}_i$.

Conditional generative model for multivariate post-processing

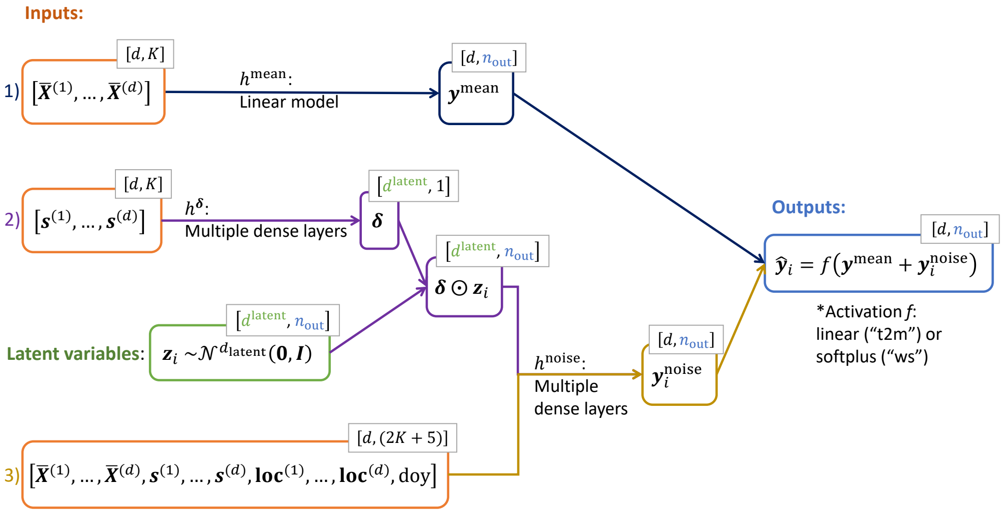

Sebastian.Lerch@kit.edu

Experimental setup

Independently for both temperature and wind speed:

- Randomly pick a station location
- Select the $d-1$ closest stations (with $d=5,10,20$)
- Generate multivariate probabilistic forecasts with benchmark and CGM approaches.

The above procedure is repeated 100 times, the results in the following summarize the variability across repetitions via boxplots.

Benchmark methods: EMOS + ECC; EMOS + GCA; DRN + ECC; DRN + GCA

Results: Univariate performance

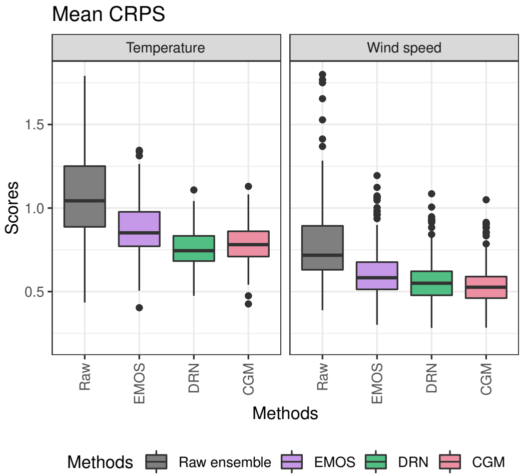

Results: Multivariate performance

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**

Results: Multivariate performance

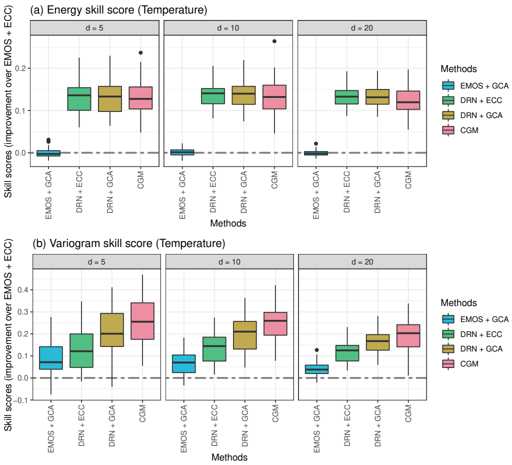

Results: Multivariate performance

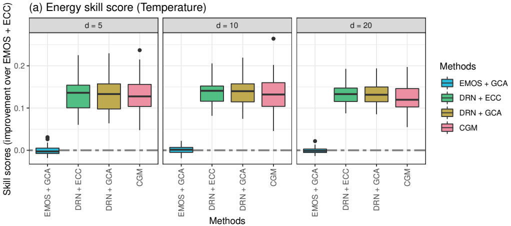

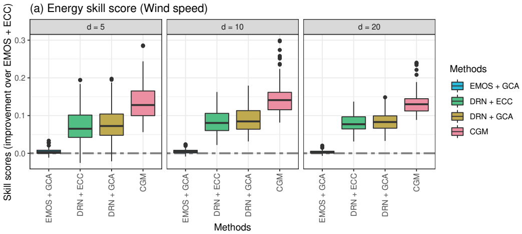

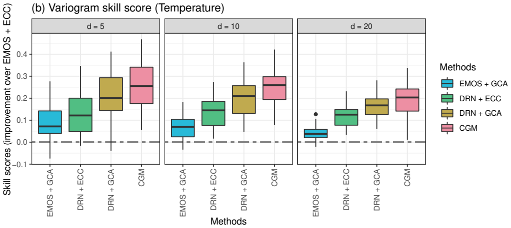

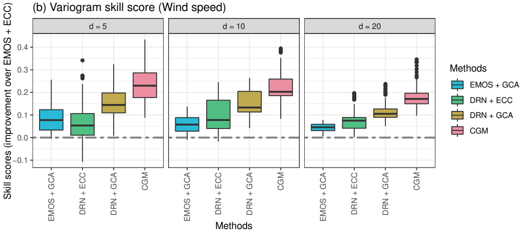

On the effects of including additional predictors

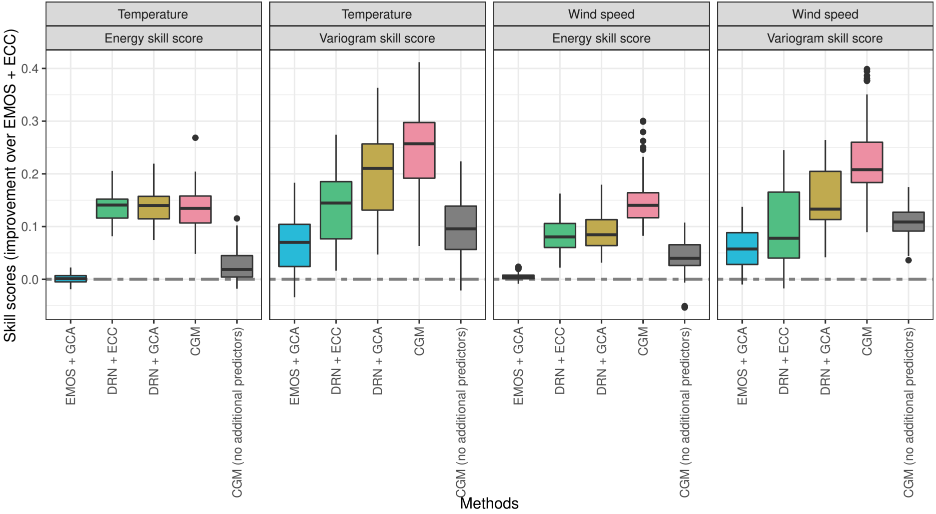

Results: Sample size

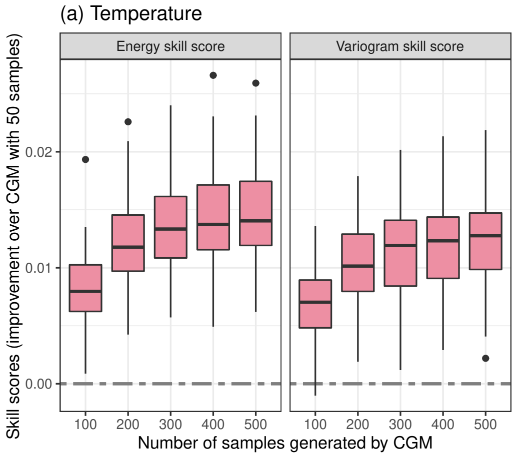

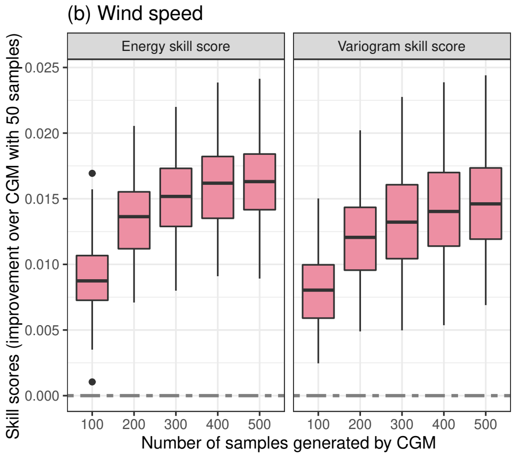

Discussion and conclusion

We propose a nonparametric multivariate post-processing method based on a conditional generative ML model.

Our generative model outperforms state-of-the-art two-step methods for multivariate post-processing in two case studies on spatial dependencies.

Discussion and conclusion

We propose a nonparametric multivariate post-processing method based on a conditional generative ML model.

Our generative model outperforms state-of-the-art two-step methods for multivariate post-processing in two case studies on spatial dependencies.

Main conceptual advantage: Ability to incorporate additional predictors.

Computational costs are manageable: even for $d = 20$, the estimation of an ensemble of 10 CGMs takes around 2 minutes on a GPU.

Sebastian.Lerch@kit.edu

Discussion and conclusion

We propose a nonparametric multivariate post-processing method based on a conditional generative ML model.

Our generative model outperforms state-of-the-art two-step methods for multivariate post-processing in two case studies on spatial dependencies.

Main conceptual advantage: Ability to incorporate additional predictors.

Computational costs are manageable: even for $d = 20$, the estimation of an ensemble of 10 CGMs takes around 2 minutes on a GPU.

Main challenges and opportunities for future work include multivariate evaluation for extremes, and applications to spatial forecast fields.

Chen, J., Janke, T., Steinke, F., and Lerch, S. (2022). **Generative machine learning methods for multivariate ensemble post-processing**, `https://doi.org/10.48550/arXiv.2211.01345`.

Code is available at `https://github.com/jieyu97/mvpp`.