

Distributional Modelling in R

Forecasting Challenge

This challenge revolves around forecasting temperatures in Graz for the next 12, 24, 36, and 48 hours. To build robust forecast models, we utilize numerical weather prediction data provided by ECMWF. While one could opt to use ECMWF forecasts directly, akin to traditional weather reports, our aim is to enhance local predictions by implementing error correction through statistical models. The typically moderate local forecast quality of ECMWF, attributed to its coarser resolution scale, underscores our main objective: developing and refining statistical models capable of effectively correcting these errors to enhance the accuracy of our temperature forecasts. This method is often referred to as "statistical post-processing" or "statistical downscaling." The observation data for 12:00:00 UTC (the response in this challenge is `obs_temp`) together with the forecast data from the ECMWF can be downloaded with

```
R> download_data <- function(data = "Graz_prepared.rds") {  
+   file <- paste0("https://nikum.org/dmr/data/", data)  
+   tdir <- tempfile()  
+   dir.create(tdir)  
+   download.file(file, file.path(tdir, data))  
+   return(readRDS(file.path(tdir, data)))  
+ }  
R> d <- download_data("Graz_prepared.rds")
```

Please note that `d` is a list containing forecasts for different forecasting horizons, with each list element represented as a "zoo" object. The corresponding data set description can be downloaded with

```
R> ds <- read.csv2("https://nikum.org/dmr/data/ECMWF_parameter_description.txt", skip = 8)
```

1. Split the data into a training set and a test set. For this purpose, designate the last four months of 2023 (September, October, November, December) as the test set, while considering all the data preceding these months as the training set.
2. Develop a comprehensive probabilistic distributional forecasting model for each forecasting horizon.
3. Additionally, visualize some forecasts to gain a better understanding of the quality of your models.
4. Using your top-performing models, calculate the mean out-of-sample log-likelihood score for the test data observations across each forecasting horizon. The model with the highest overall mean score will be declared the winner of the challenge!