

Probabilistic Weather Forecasting in R

by Chris Fraley, Adrian Raftery, Tilmann Gneiting, McLean Sloughter and Veronica Berrocal

Abstract This article describes two R packages for probabilistic weather forecasting, **ensembleBMA**, which offers ensemble postprocessing via Bayesian model averaging (BMA), and **ProbForecastGOP**, which implements the geostatistical output perturbation (GOP) method. BMA forecasting models use mixture distributions, in which each component corresponds to an ensemble member, and the form of the component distribution depends on the weather parameter (temperature, quantitative precipitation or wind speed). The model parameters are estimated from training data. The GOP technique uses geostatistical methods to produce probabilistic forecasts of entire weather fields for temperature or pressure, based on a single numerical forecast on a spatial grid. Both packages include functions for evaluating predictive performance, in addition to model fitting and forecasting.

Introduction

Over the past two decades, weather forecasting has experienced a paradigm shift towards probabilistic forecasts, which take the form of probability distributions over future weather quantities and events. Probabilistic forecasts allow for optimal decision making for many purposes, including air traffic control, ship routing, agriculture, electricity generation and weather-risk finance.

Up to the early 1990s, most weather forecasting was deterministic, meaning that only one “best” forecast was produced by a numerical model. The recent advent of ensemble prediction systems marks a radical change. An ensemble forecast consists of multiple numerical forecasts, each computed in a different way. Statistical postprocessing is then used to convert the ensemble into calibrated and sharp probabilistic forecasts (Gneiting and Raftery, 2005).

The ensembleBMA package

The **ensembleBMA** package (Fraley et al., 2010) offers statistical postprocessing of forecast ensembles via Bayesian model averaging (BMA). It provides functions for model fitting and forecasting with ensemble data that may include missing and/or exchangeable members. The modeling functions estimate BMA parameters from training data via the EM algorithm. Currently available options are normal mixture models (appropriate for temperature or

pressure), mixtures of gamma distributions (appropriate for wind speed), and Bernoulli-gamma mixtures with a point mass at zero (appropriate for quantitative precipitation). Functions for verification that assess predictive performance are also available.

The BMA approach to the postprocessing of ensemble forecasts was introduced by Raftery et al. (2005) and has been developed in Berrocal et al. (2007), Sloughter et al. (2007), Wilson et al. (2007), Fraley et al. (2010) and Sloughter et al. (2010). Detail on verification procedures can be found in Gneiting and Raftery (2007) and Gneiting et al. (2007).

"ensembleData" objects

Ensemble forecasting data for weather typically includes some or all of the following information:

- ensemble member forecasts
- initial date
- valid date
- forecast hour (prediction horizon)
- location (latitude, longitude, elevation)
- station and network identification

The *initial date* is the day and time at which initial conditions are provided to the numerical weather prediction model, to run forward the partial differential equations that produce the members of the forecast ensemble. The *forecast hour* is the prediction horizon or time between initial and valid dates. The ensemble member forecasts then are valid for the hour and day that correspond to the forecast hour ahead of the initial date. In all the examples and illustrations in this article, the prediction horizon is 48 hours.

For use with the **ensembleBMA** package, data must be organized into an "ensembleData" object that minimally includes the ensemble member forecasts. For model fitting and verification, the corresponding weather observations are also needed. Several of the model fitting functions can produce forecasting models over a sequence of dates, provided that the "ensembleData" are for a single prediction horizon. Attributes such as station and network identification, and latitude and longitude, may be useful for plotting and/or analysis but are not currently used in any of the modeling functions. The "ensembleData" object facilitates preservation of the data as a unit for use in processing by the package functions.

Here we illustrate the creation of an "ensembleData" object called `srftData` that corresponds to the `srft` data set of surface temperature (Berrocal et al., 2007):

```
data(srft)
members <- c("CMCG", "ETA", "GASP", "GFS",
             "JMA", "NGPS", "TCWB", "UKMO")
srftData <-
  ensembleData(forecasts = srft[,members],
              dates = srft$date,
              observations = srft$obs,
              latitude = srft$lat,
              longitude = srft$lon,
              forecastHour = 48)
```

The dates specification in an "ensembleData" object refers to the valid dates for the forecasts.

Specifying exchangeable members. Forecast ensembles may contain members that can be considered exchangeable (arising from the same generating mechanism, such as random perturbations of a given set of initial conditions) and for which the BMA parameters, such as weights and bias correction coefficients, should be the same. In **ensembleBMA**, exchangeability is specified by supplying a vector that represents the grouping of the ensemble members. The non-exchangeable groups consist of singleton members, while exchangeable members belong to the same group. See Fraley et al. (2010) for a detailed discussion.

Specifying dates. Functions that rely on the **chron** package (James and Hornik, 2010) are provided for converting to and from Julian dates. These functions check for proper format ('YYYYMMDD' or 'YYYYMMDDHH').

BMA forecasting

BMA generates full predictive probability density functions (PDFs) for future weather quantities. Examples of BMA predictive PDFs for temperature and precipitation are shown in Figure 1.

Surface temperature example. As an example, we fit a BMA normal mixture model for forecasts of surface temperature valid January 31, 2004, using the **srft** training data. The "ensembleData" object **srftData** created in the previous section is used to fit the predictive model, with a rolling training period of 25 days, excluding the two most recent days because of the 48 hour prediction horizon.

One of several options is to use the function **ensembleBMA** with the valid date(s) of interest as input to obtain the associated BMA fit(s):

```
srftFit <-
  ensembleBMA(srftData, dates = "2004013100",
             model = "normal", trainingDays = 25)
```

¹The package implements the original BMA method of Raftery et al. (2005) and Slougher et al. (2007), in which there is a single, constant bias correction term over the whole domain. Model biases are likely to differ by location, and there are newer methods that account for this (Gel, 2007; Mass et al., 2008; Kleiber et al., in press).

When no dates are specified, a model fit is produced for each date for which there are sufficient training data for the desired rolling training period.

The BMA predictive PDFs can be plotted as follows, with Figure 1 showing an example:

```
plot(srftFit, srftData, dates = "2004013100")
```

This steps through each location on the given dates, plotting the corresponding BMA PDFs.

Alternatively, the modeling process for a single date can be separated into two steps: first extracting the training data, and then fitting the model directly using the **fitBMA** function. See Fraley et al. (2007) for an example. A limitation of this approach is that date information is not automatically retained.

Forecasting is often done on grids that cover an area of interest, rather than at station locations. The dataset **srftGrid** provides ensemble forecasts of surface temperature initialized on January 29, 2004 and valid for January 31, 2004 at grid locations in the same region as that of the **srft** stations.

Quantiles of the BMA predictive PDFs at the grid locations can be obtained with the function **quantileForecast**:

```
srftGridForc <- quantileForecast(srftFit,
                               srftGridData, quantiles = c(.1, .5, .9))
```

Here **srftGridData** is an "ensembleData" object created from **srftGrid**, and **srftFit** provides a forecasting model for the corresponding date.¹ The forecast probability of temperatures below freezing at the grid locations can be computed with the **cdf** function, which evaluates the BMA cumulative distribution function:

```
probFreeze <- cdf(srftFit, srftGridData,
                 date = "2004013100", value = 273.15)
```

In the **srft** and **srftGrid** datasets, temperature is recorded in degrees Kelvin (K), so freezing occurs at 273.15 K.

These results can be displayed as image plots using the **plotProbcast** function, as shown in Figure 2, in which darker shades represent higher probabilities. The plots are made by binning values onto a plotting grid, which is the default in **plotProbcast**. Loading the **fields** (Furrer et al., 2010) and **maps** (Brownrigg and Minka, 2011) packages enables display of the country and state outlines, as well as a legend.

Precipitation example. The **propFit** dataset consists of the fitted BMA parameters for 48 hour ahead forecasts of daily precipitation accumulation (in hundredths of an inch) over the U.S. Pacific Northwest from December 12, 2002 through March 31, 2005, as described by Slougher et al. (2007). The fitted models are Bernoulli-gamma mixtures with a point




Figure 1: BMA predictive distributions for temperature (in degrees Kelvin) valid January 31, 2004 (left) and for precipitation (in hundredths of an inch) valid January 15, 2003, at Port Angeles, Washington at 4PM local time, based on the eight-member University of Washington Mesoscale Ensemble (Grimt and Mass, 2002; Eckel and Mass, 2005). The thick black curve is the BMA PDF, while the colored curves are the weighted PDFs of the constituent ensemble members. The thin vertical black line is the median of the BMA PDF (occurs at or near the mode in the temperature plot), and the dashed vertical lines represent the 10th and 90th percentiles. The orange vertical line is at the verifying observation. In the precipitation plot (right), the thick vertical black line at zero shows the point mass probability of no precipitation (47%). The densities for positive precipitation amounts have been rescaled, so that the maximum of the thick black BMA PDF agrees with the probability of precipitation (53%).

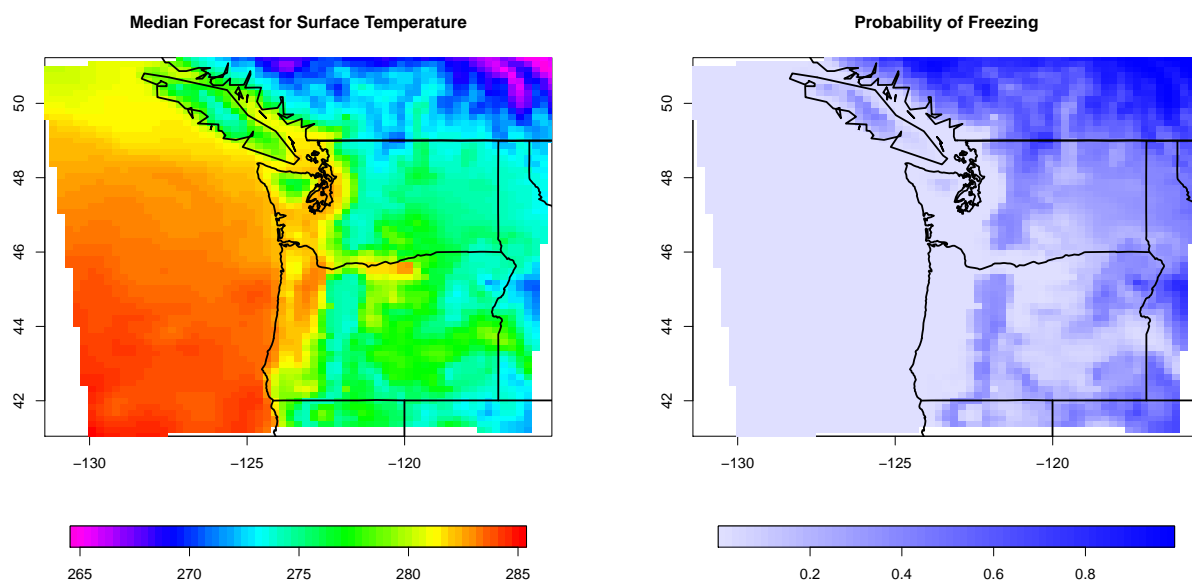


Figure 2: Image plots of the BMA median forecast for surface temperature and BMA probability of freezing over the Pacific Northwest, valid January 31, 2004.

mass at zero that apply to the cube root transformation of the ensemble forecasts and observed data. A rolling training period of 30 days is used. The dataset used to obtain `prcpFit` is not included in the package on account of its size. However, the corresponding "ensembleData" object can be constructed in the same way as illustrated for the surface temperature data, and the modeling process also is analogous, except that the "gamma0" model for quantitative precipitation is used in lieu of the "normal" model.

The `prcpGrid` dataset contains gridded ensemble forecasts of daily precipitation accumulation in the same region as that of `prcpFit`, initialized January 13, 2003 and valid January 15, 2003. The BMA median and upper bound (90th percentile) forecasts can be obtained and plotted as follows:

```
data(prcpFit)

prcpGridForc <- quantileForecast(
  prcpFit, prcpGridData, date = "20030115",
  q = c(0.5, 0.9))
```

Here `prcpGridData` is an "ensembleData" object created from the `prcpGrid` dataset. The 90th percentile plot is shown in Figure 3. The probability of precipitation and the probability of precipitation above .25 inches can be obtained as follows:

```
probPrecip <- 1 - cdf(prcpFit, prcpGridData,
  date = "20030115", values = c(0, 25))
```

The plot for the BMA forecast probability of precipitation accumulation exceeding .25 inches is also shown in Figure 3.

Verification

The **ensembleBMA** functions for verification can be used whenever observed weather conditions are available. Included are functions to compute verification rank and probability integral transform histograms, the mean absolute error, continuous ranked probability score, and Brier score.

Mean absolute error, continuous ranked probability score, and Brier score. In the previous section, we obtained a gridded BMA forecast of surface temperature valid January 31, 2004 from the `srft` data set. To obtain forecasts at station locations, we apply the function `quantileForecast` to the model fit `srftFit`:

```
srftForc <- quantileForecast(srftFit,
  srftData, quantiles = c(.1, .5, .9))
```

The BMA quantile forecasts can be plotted together with the observed weather conditions using the function `plotProbcast` as shown in Figure 4. Here the R core function `loess` was used to interpolate from the

station locations to a grid for surface plotting. It is also possible to request image or perspective plots, or contour plots.

The mean absolute error (MAE) and mean continuous ranked probability score (CRPS; e.g., Gneiting and Raftery, 2007) can be computed with the functions `CRPS` and `MAE`:

```
CRPS(srftFit, srftData)
# ensemble      BMA
# 1.945544 1.490496

MAE(srftFit, srftData)
# ensemble      BMA
# 2.164045 2.042603
```

The function `MAE` computes the mean absolute difference between the ensemble or BMA median forecast² and the observation. The BMA CRPS is obtained via Monte Carlo simulation and thus may vary slightly in replications. Here we compute these measures from forecasts valid on a single date; more typically, the CRPS and MAE would be computed from a range of dates and the corresponding predictive models.

Brier scores (see e.g., Jolliffe and Stephenson, 2003; Gneiting and Raftery, 2007) for probability forecasts of the binary event of exceeding an arbitrary precipitation threshold can be computed with the function `brierScore`.

Assessing calibration. Calibration refers to the statistical consistency between the predictive distributions and the observations (Gneiting et al., 2007). The *verification rank histogram* is used to assess calibration for an ensemble forecast, while the *probability integral transform* (PIT) histogram assesses calibration for predictive PDFs, such as the BMA forecast distributions.

The verification rank histogram plots the rank of each observation relative to the combined set of the ensemble members and the observation. Thus, it equals one plus the number of ensemble members that are smaller than the observation. The histogram allows for the visual assessment of the calibration of the ensemble forecast (Hamill, 2001). If the observation and the ensemble members are exchangeable, all ranks are equally likely, and so deviations from uniformity suggest departures from calibration. We illustrate this with the `srft` dataset, starting at January 30, 2004:

```
use <- ensembleValidDates(srftData) >=
  "2004013000"

srftVerifRank <- verifRank(
  ensembleForecasts(srftData[use,]),
  ensembleVerifObs(srftData[use,]))
```

²Raftery et al. (2005) employ the BMA predictive mean rather than the predictive median.

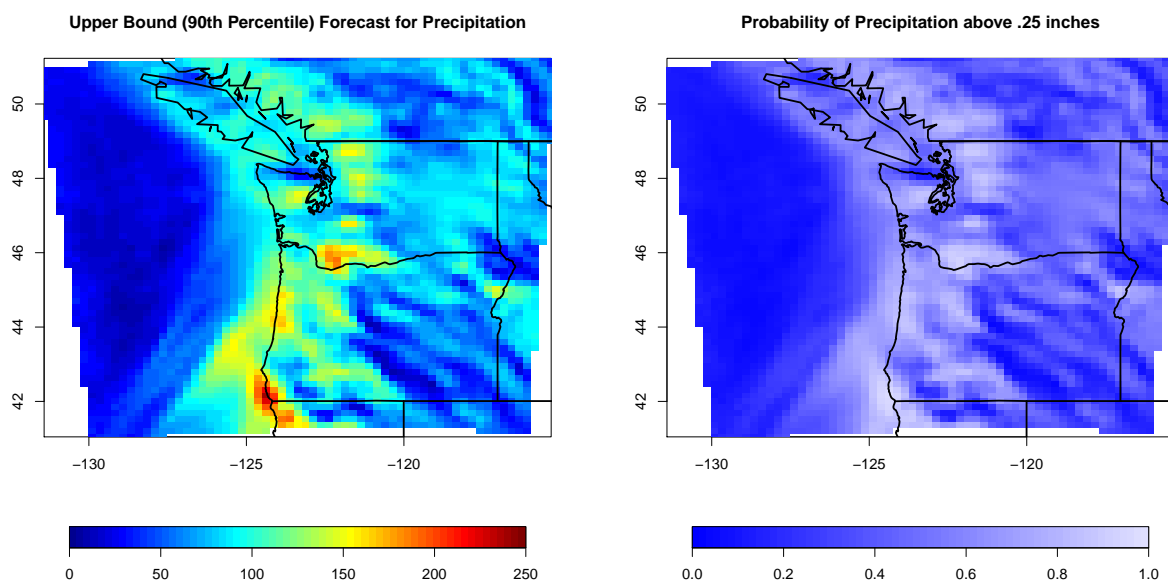


Figure 3: Image plots of the BMA upper bound (90th percentile) forecast of precipitation accumulation (in hundredths of an inch), and the BMA probability of precipitation exceeding .25 inches over the Pacific Northwest, valid January 15, 2003.

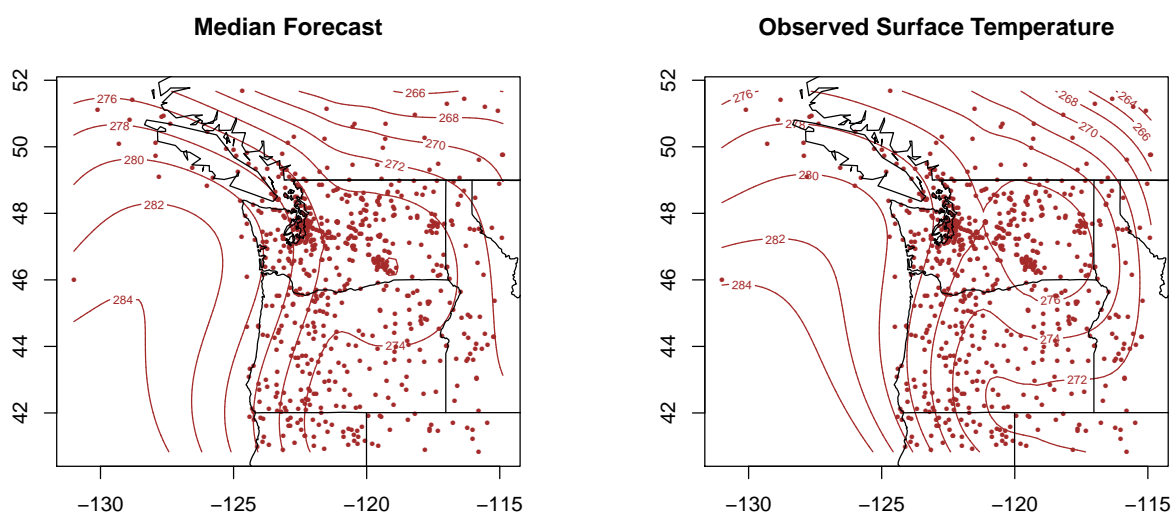


Figure 4: Contour plots of the BMA median forecast of surface temperature and verifying observations at station locations in the Pacific Northwest, valid January 31, 2004 (srft dataset). The plots use loess fits to the forecasts and observations at the station locations, which are interpolated to a plotting grid. The dots represent the 715 observation sites.

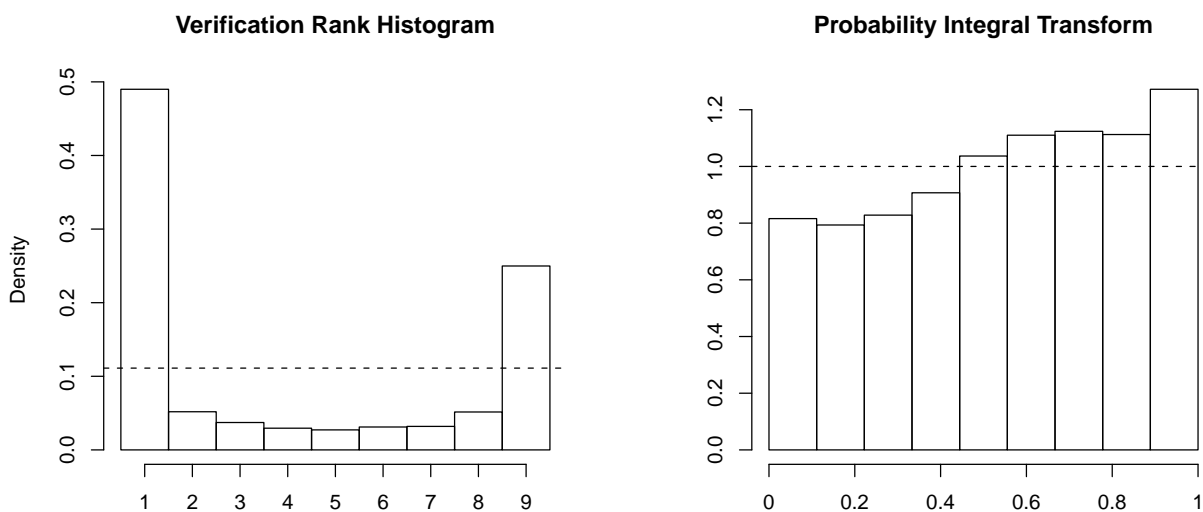


Figure 5: Verification rank histogram for ensemble forecasts, and PIT histogram for BMA forecast PDFs for surface temperature over the Pacific Northwest in the `srft` dataset valid from January 30, 2004 to February 28, 2004. More uniform histograms correspond to better calibration.

```
k <- ensembleSize(srftData)

hist(srftVerifRank, breaks = 0:(k+1),
     prob = TRUE, xaxt = "n", xlab = "",
     main = "Verification Rank Histogram")
axis(1, at = seq(.5, to = k+.5, by = 1),
     labels = 1:(k+1))
abline(h=1/(ensembleSize(srftData)+1), lty=2)
```

The resulting rank histogram composites ranks spatially and is shown in Figure 5. The U shape indicates a lack of calibration, in that the ensemble forecast is underdispersed.

The PIT is the value that the predictive cumulative distribution function attains at the observation, and is a continuous analog of the verification rank. The function `pit` computes it. The PIT histogram allows for the visual assessment of calibration and is interpreted in the same way as the verification rank histogram. We illustrate this on BMA forecasts of surface temperature obtained for the entire `srft` data set using a 25 day training period (forecasts begin on January 30, 2004 and end on February 28, 2004):

```
srftFitALL <- ensembleBMA(srftData,
                        trainingDays = 25)

srftPIT <- pit(srftFitALL, srftData)

hist(srftPIT, breaks = (0:(k+1))/(k+1),
     xlab="", xaxt="n", prob = TRUE,
     main = "Probability Integral Transform")
axis(1, at = seq(0, to = 1, by = .2),
     labels = (0:5)/5)
abline(h = 1, lty = 2)
```

The resulting PIT histogram is shown in Figure 5. It shows signs of negative bias, which is not surprising because it is based on only about a month of verifying data. We generally recommend computing the PIT histogram for longer periods, ideally at least a year, to avoid its being dominated by short-term and seasonal effects.

The ProbForecastGOP package

The **ProbForecastGOP** package (Berrocal et al., 2010) generates probabilistic forecasts of entire weather fields using the geostatistical output perturbation (GOP) method of Gel et al. (2004). The package contains functions for the GOP method, a wrapper function named `ProbForecastGOP`, and a plotting utility named `plotfields`. More detailed information can be found in the PDF document installed in the package directory at `'ProbForecastGOP/docs/vignette.pdf'` or in the help files.

The GOP method uses geostatistical methods to produce probabilistic forecasts of entire weather fields for temperature and pressure, based on a single numerical weather forecast on a spatial grid. The method involves three steps:

- an initial step in which linear regression is used for bias correction, and the empirical variogram of the forecast errors is computed;
- an estimation step in which the weighted least squares method is used to fit a parametric model to the empirical variogram; and
- a forecasting step in which a statistical ensemble of weather field forecasts is generated,

by simulating realizations of the error field from the fitted geostatistical model, and adding them to the bias-corrected numerical forecast field.

Empirical variogram

In the first step of the GOP method, the empirical variogram of the forecast errors is found. Variograms are used in spatial statistics to characterize variability in spatial data. The empirical variogram plots one-half the mean squared difference between paired observations against the distance separating the corresponding sites. Distances are usually binned into intervals, whose midpoints are used to represent classes.

In `ProbForecastGOP`, four functions compute empirical variograms. Two of them, `avg.variog` and `avg.variog.dir`, compute forecast errors over temporal replications, while the other two, `Emp.variog` and `EmpDir.variog`, average forecast errors temporally, and only then compute variograms. Alternatively, one can use the wrapper function `ProbForecastGOP` with argument `out = "VARIOG"`.

Parameter estimation

The second step in the GOP method consists of fitting a parametric model to the empirical variogram of the forecast errors. This is done by the `Variog.fit` function using the weighted least squares approach. Alternatively, the wrapper function `ProbForecastGOP` with entry `out` set equal to "FIT" can be employed. Figure 6 illustrates the result for forecasts of surface temperature over the Pacific Northwest in January through June 2000.




Figure 6: Empirical variogram of temperature forecast errors and fitted exponential variogram model.

The parametric models implemented in `Variog.fit` are the exponential, spherical, Gaussian, generalized Cauchy and Whittle-Matérn, with further detail available in the help files. The function `linesmodel` computes these parametric models.

Generating ensemble members

The final step in the GOP method involves generating multiple realizations of the forecast weather field. Each realization provides a member of a statistical ensemble of weather field forecasts, and is obtained by simulating a sample path of the fitted error field, and adding it to the bias-adjusted numerical forecast field.

This is achieved by the function `Field.sim`, or by the wrapper function `ProbForecastGOP` with the entry `out` set equal to "SIM". Both options depend on the `GaussRF` function in the `RandomFields` package that simulates Gaussian random fields (Schlather, 2011).

The output of the functions is both numerical and graphical, in that they return members of the forecast field ensemble, quantile forecasts at individual locations, and plots of the ensemble members and quantile forecasts. The plots are created using the `image.plot`, `US` and `world` functions in the `fields` package (Furrer et al., 2010). As an illustration, Figure 7 shows a member of a forecast field ensemble of surface temperature over the Pacific Northwest valid January 12, 2002.




Figure 7: A member of a 48-hour ahead forecast field ensemble of surface temperature (in degrees Celsius) over the Pacific Northwest valid January 12, 2002 at 4PM local time.

Summary

We have described two packages, **ensembleBMA** and **ProbForecastGOP**, for probabilistic weather forecasting. Both packages provide functionality to fit forecasting models to training data.

In **ensembleBMA**, parametric mixture models, in which the components correspond to the members of an ensemble, are fit to a training set of ensemble forecasts, in a form that depends on the weather parameter. These models are then used to postprocess ensemble forecasts at station or grid locations.

In **ProbForecastGOP**, a parametric model is fit to the empirical variogram of the forecast errors from a single member, and forecasts are obtained by simulating realizations of the fitted error field and adding them to the bias-adjusted numerical forecast field. The resulting probabilistic forecast is for an entire weather field, and shows physically realistic spatial features.

Supplementing model fitting and forecasting, both packages provide functionality for verification, allowing the quality of the forecasts produced to be assessed.

Acknowledgements

This research was supported by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the Office of Naval Research under Grant N00014-01-10745, by the National Science Foundation under grants ATM-0724721 and DMS-0706745, and by the Joint Ensemble Forecasting System (JEFS) under subcontract S06-47225 from the University Corporation for Atmospheric Research (UCAR). We are indebted to Cliff Mass and Jeff Baars of the University of Washington Department of Atmospheric Sciences for sharing insights, expertise and data, and thank Michael Scheuerer for comments. This paper also benefitted from the critiques of two anonymous reviewers and the associate editor.

Bibliography

- V. J. Berrocal, Y. Gel, A. E. Raftery, and T. Gneiting. *ProbForecastGOP: Probabilistic weather field forecast using the GOP method*, 2010. URL <http://CRAN.R-project.org/package=ProbForecastGOP>. R package version 1.3.2.
- V. J. Berrocal, A. E. Raftery, and T. Gneiting. Combining spatial statistical and ensemble information in probabilistic weather forecasts. *Monthly Weather Review*, 135:1386–1402, 2007.
- R. Brownrigg and T. P. Minka. *maps: Draw geographical maps*, 2011. URL <http://CRAN.R-project.org/package=maps>. R package version 2.1-6. S original by Richard A. Becker and Allan R. Wilks, R port by Ray Brownrigg, enhancements by Thomas P. Minka.
- F. A. Eckel and C. F. Mass. Effective mesoscale, short-range ensemble forecasting. *Weather and Forecasting*, 20:328–350, 2005.
- C. Fraley, A. E. Raftery, T. Gneiting, and J. M. Sloughter. *ensembleBMA: Probabilistic forecasting using ensembles and Bayesian Model Averaging*, 2010. URL <http://CRAN.R-project.org/package=ensembleBMA>. R package version 4.5.
- C. Fraley, A. E. Raftery, T. Gneiting, and J. M. Sloughter. *ensembleBMA: An R package for probabilistic forecasting using ensembles and Bayesian model averaging*. Technical Report 516, University of Washington, Department of Statistics, 2007. (revised 2011).
- C. Fraley, A. E. Raftery, and T. Gneiting. Calibrating multi-model forecasting ensembles with exchangeable and missing members using Bayesian model averaging. *Monthly Weather Review*, 138:190–202, 2010.
- R. Furrer, D. Nychka, and S. Sain. *fields: Tools for spatial data*, 2010. URL <http://CRAN.R-project.org/package=fields>. R package version 6.3.
- Y. Gel, A. E. Raftery, and T. Gneiting. Calibrated probabilistic mesoscale weather field forecasting: the Geostatistical Output Perturbation (GOP) method (with discussion). *Journal of the American Statistical Association*, 99:575–590, 2004.
- Y. R. Gel. Comparative analysis of the local observation-based (LOB) method and the non-parametric regression-based method for gridded bias correction in mesoscale weather forecasting. *Weather and Forecasting*, 22:1243–1256, 2007.
- T. Gneiting and A. E. Raftery. Weather forecasting with ensemble methods. *Science*, 310:248–249, 2005.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.
- T. Gneiting, F. Balabdaoui, and A. E. Raftery. Probabilistic forecasts, calibration, and sharpness. *Journal of the Royal Statistical Society, Series B*, 69:243–268, 2007.
- E. P. Grimit and C. F. Mass. Initial results for a mesoscale short-range ensemble forecasting system over the Pacific Northwest. *Weather and Forecasting*, 17:192–205, 2002.
- T. M. Hamill. Interpretation of rank histograms for verifying ensemble forecasts. *Monthly Weather Review*, 129:550–560, 2001.

- D. James and K. Hornik. *chron: Chronological objects which can handle dates and times*, 2010. URL <http://CRAN.R-project.org/package=chron>. R package version 2.3-39. S original by David James, R port by Kurt Hornik.
- I. T. Jolliffe and D. B. Stephenson, editors. *Forecast Verification: A Practitioner's Guide in Atmospheric Science*. Wiley, 2003.
- W. Kleiber, A. E. Raftery, J. Baars, T. Gneiting, C. Mass, and E. P. Gritmit. Locally calibrated probabilistic temperature forecasting using geostatistical model averaging and local Bayesian model averaging. *Monthly Weather Review*, 2011 (in press).
- C. F. Mass, J. Baars, G. Wedam, E. P. Gritmit, and R. Steed. Removal of systematic model bias on a model grid. *Weather and Forecasting*, 23:438–459, 2008.
- A. E. Raftery, T. Gneiting, F. Balabdaoui, and M. Polakowski. Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133:1155–1174, 2005.
- M. Schlather. *RandomFields: Simulation and analysis tools for random fields*, 2011. URL <http://CRAN.R-project.org/package=RandomFields>. R package version 2.0.45.
- J. M. Sloughter, A. E. Raftery, T. Gneiting, and C. Fraley. Probabilistic quantitative precipitation forecasting using Bayesian model averaging. *Monthly Weather Review*, 135:3209–3220, 2007.
- J. M. Sloughter, T. Gneiting, and A. E. Raftery. Probabilistic wind speed forecasting using ensembles and Bayesian model averaging. *Journal of the American Statistical Association*, 105:25–35, 2010.
- L. J. Wilson, S. Beauregard, A. E. Raftery, and R. Verret. Calibrated surface temperature forecasts from the Canadian ensemble prediction system using Bayesian model averaging. *Monthly Weather Review*, 135:1364–1385, 2007.

Chris Fraley
Department of Statistics
University of Washington
Seattle, WA 98195-4322 USA
fraley@stat.washington.edu

Adrian Raftery
Department of Statistics
University of Washington
Seattle, WA 98195-4322 USA
raftery@u.washington.edu

Tilmann Gneiting
Institute of Applied Mathematics
Heidelberg University
69120 Heidelberg, Germany
t.gneiting@uni-heidelberg.de

McLean Sloughter
Department of Mathematics
Seattle University
Seattle, WA 98122 USA
sloughtj@seattleu.edu

Veronica Berrocal
School of Public Health
University of Michigan
1415 Washington Heights
Ann Arbor, MI 48109-2029 USA
berrocal@umich.edu