`ensembleBMA`: An `R` Package for Probabilistic Forecasting using Ensembles and Bayesian Model Averaging *

Chris Fraley, Adrian E. Raftery
Tilmann Gneiting, J. McLean Sloughter

Technical Report No. 516

Department of Statistics
University of Washington
Box 354322
Seattle, WA 98195-4322 USA

first version August 15, 2007; this version February 19, 2013

Abstract

`ensembleBMA` is a contributed `R` package for probabilistic forecasting using ensemble post-processing via Bayesian Model Averaging. It provides functions for modeling and forecasting with data that may include missing ensemble member forecasts. The modeling can also account for exchangeable ensemble members. The modeling functions estimate model parameters from training data via the EM algorithm for normal mixture models (appropriate for temperature or pressure), mixtures of gamma distributions (appropriate for maximum wind speed), and mixtures of gamma distributions with a point mass at 0 (appropriate for quantitative precipitation). Also included are functions for forecasting from these models, as well as functions for verification to assess forecasting performance.

*Thanks go to Veronica Berrocal and Patrick Tewson for lending their expertise on a number of important issues, to Michael Polakowski for his work on an earlier version of the package, to Thordis Thorarinsdottir and Bobby Yuen for complementary work on ensemble MOS, and to a number of users whose observations and suggestions have helped improve this package. We are also indebted to Cliff Mass, Jeff Baars, and Eric Grimit for many helpful discussions on forecasting and verification, and for sharing data. Supported by the DoD Multidisciplinary Research Initiative (MURI) administered by the Office of Naval Research under grant N00014-01-10745 and by the National Science Foundation under grants ATM-0724721, DMS-0706745, and subcontract S06-47225 to the University Corporation for Atmospheric Research (UCAR). The latter project is part of the Joint Ensemble Forecast System (JEFS) funded by the Air Force Weather Agency through NSF.

Contents

List of Tables

List of Figures

1 Overview

This document describes the `ensembleBMA` package for probabilistic forecasting using ensemble postprocessing via Bayesian Model Averaging (BMA), written in the `R` language. The package offers the following capabilities:

- Fitting BMA models to ensemble forecasting data with verifying observations. Modeling options are as follows:

  - mixtures of normals for temperature and pressure
  - mixtures of gammas for maximum wind speed
  - mixtures of gammas with a point mass at 0 for quantitative precipitation

  The modeling can accomodate exchangeable ensemble members, as well as missing member forecasts (Fraley et al. 2010).

- Producing quantile forecasts from fitted BMA models. Forecasting from data with missing ensemble members is possible.

- Computing continuous ranked probability scores, Brier scores, and other measures for assessing BMA forecasting performance.

- Displaying forecast and verification results.

The modeling methodology used in `ensembleBMA` was introduced in Raftery et al. (2005). More detail on the models and verification procedures can be found in Gneiting et al. (2007), Gneiting and Raftery (2007), Sloughter et al. (2007), Wilson et al. (2007), Sloughter et al. (2009) and Fraley et al. (2010). These methods are based on forecast ensembles; for an overview of ensemble weather forecasting, see Gneiting and Raftery (2005).

To use the `ensembleBMA` package, download it from the Comprehensive R Archive Network (CRAN) `http://cran.r-project.org`. Follow the instructions for installing `R` packages on your machine, and then do

```
> library(ensembleBMA)
```

inside `R` in order to use the software. Throughout this document it will be assumed that these steps have been taken before running the examples.

2 `ensembleData` objects

Modeling and forecasting functions in the `ensembleBMA` package require that the data be organized into an `ensembleData` object that includes the ensemble forecasts with their valid dates. Observed weather conditions are also needed for modeling and verification. Other attributes such as latitude and longitude, station and network identification, and elevation may be useful for plotting and/or analysis. For the batch-processing modeling functions that underly (the function) `ensembleBMA`, the data in an `ensembleData` object are expected to be for a single forecast hour and initialization time. The forecast hour (the time interval between the initialization time and the forecast time) must be specified (in units of hours) for data processed by these functions in order to determine the appropriate training lag. The initialization time may also be specified to help ensure that models and data are consistent within an analysis. The `ensembleData` object facilitates preservation of the data as a unit for use in processing by the package functions.

As an example, we create an `ensembleData` object called `srftData` corresponding to the `srft` data set of surface temperatures (see Berrocal et al 2010):

```
data(srft)

memberLabels <- c("CMCG","ETA","GASP","GFS","JMA","NGPS","TCWB","UKMO")

srftData <- ensembleData(forecasts = srft[,memberLabels],
                dates = srft$date, observations = srft$obs,
                latitude = srft$lat, longitude = srft$lon,
                forecastHour = 48, initializationTime = "00")
```

The labels of the member forecasts should be consistent for `ensembleData` objects used within an analysis, because they are used to match member names in data with the BMA model weights and parameters for forecasting and verification for consistency in composition and order among datasets.

Specifying dates. When dates are included in an `ensembleData` object, they must be specified as a character vector (or its factor equivalent) using strings in the form `YYYYMMDDHH` or `YYYYMMDD`, in which `YYYY` specifies the year, `MM` the month, `DD` the day, and (optionally) `HH` the hour. The `ensembleData` function checks the date format for correctness. A function `julTOymdh` is provided for converting vectors of Julian dates to vectors of equivalent dates in the required format, along with another function `ymdhTOjul` that does the reverse. These functions rely on the `chron` package (Hornik 1999).

Specifying exchangeable ensemble members. Forecast ensembles may contain members that can be considered exchangeable or interchangeable; that is, their forecasts can be assumed to come from the same distribution. In such cases, parameters in the BMA model (including weights and bias correction coefficients) should be constrained to be the same among exchangeable members. In `ensembleBMA`, exchangeability is specified by supplying a vector representing a grouping of the ensemble members in the `exchangeable` argument when creating `ensembleData` objects. The non-interchangeable groups consist of singleton members, while exchangeable members would belong to the same group. As an illustration, suppose the `ETA` and `GFS` members are exchangeable in the example above, but all other members are non-interchangeable. The corresponding `ensembleData` object could be created as follows:

```
data(srft)

memberLabels <- c("CMCG","ETA","GASP","GFS","JMA","NGPS","TCWB","UKMO")
```

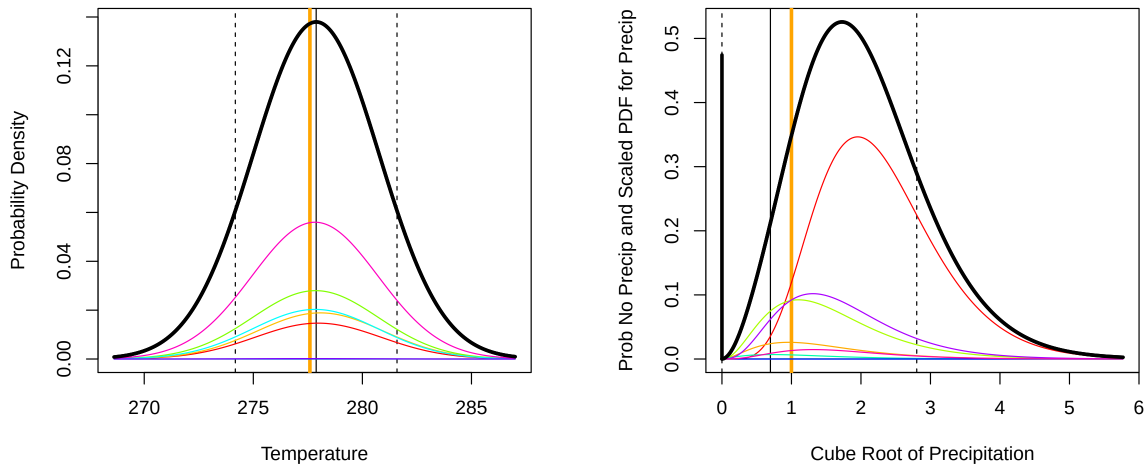

Figure 1: BMA predictive distributions for temperature (in degrees Kelvin) valid January 31, 2004 (left) and for precipitation (in hundredths of an inch) valid January 15, 2003 (right), at Port Angeles, Washington at 4PM local time, based on the eight-member University of Washington Mesoscale Ensemble (Grimit 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%).

```
exGroups <- c( CMCG=1, ETA=2, GASP=3, GFS=2, JMA=4, NGPS=5, TCWB=6, UKMO=7)

srftDataX <- ensembleData(forecasts = srft[,memberLabels],
                  dates = srft$date, observations = srft$obs,
                  latitude = srft$lat, longitude = srft$lon,
                  forecastHour = 48, initializationTime = "00",
                  exchangeable = exGroups)
```

The weights and parameters in a BMA model fit to `srtfDataX` will be equal for the `ETA` and `GFS` members.

See Fraley et al. 2010 for a detailed discussion of how exchangeability is handled in BMA postprocessing.

3 BMA Forecasting

BMA produces a probability distribution function (PDF) for the weather data, from which forecasts can be made a verified. Examples of BMA predictive distributions for temperature and precipitation are shown in Figure 1.

3.1 Surface Temperature Example.

As an example, we model 48-hour surface temperature for January 31, 2004 from ensemble forecasts and observations at station locations as given in the `srft` data set provided in the `ensembleBMA` package. The model fits a mixture of normals to the ensemble forecasts and observed data. We use `srftData`, one of the `ensembleData` objects created in the previous section, in the modeling. A training period of 25 days is used, with a lag of 2 days in the training data since the `srft` dataset is for forecast hour 48. The data is fitted with a mixture of normals as appropriatpe for temperature.

There are several options for obtaining the model. One is to use the function `ensembleBMA` with the valid date (or dates) of interest as input to obtain the associated BMA model (or models).

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

It should be noted that the `ensembleBMA` function will produce a model for each valid date specified in the `dates` argument, provided that the date is consistent with the available data and the training period. When no dates are specified, the `ensembleBMA` function will produce a model for each date for which there is sufficient training data in the input data for the desired training period. The result of applying `ensembleBMA` with multiple dates can be used to obtain forecasting models for each of those dates. The BMA predictive distributions can be plotted (as in Figure 1) as follows:

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

This steps through each location on each date, plotting the corresponding BMA PDF.

The modeling process for a single date can also be separated into two steps: extraction of the training data for the desired date, and then fitting the model directly with `fitBMA`.

```
train <- trainingData( srftData, date = "2004013100",
                               trainingDays = 25)
srftTrainFit <- fitBMA( train,  model = "normal")
```

A limitation of the two-step process is that date information is not retained as part of the model.

Forecasting is typically done on grids covering an area of interest rather than at station locations. The dataset `srftGrid` included in the `ensembleBMA` package gives forecasts of surface temperature initialized on January 29, 2004 and valid for January 31, 2004 at grid locations in the region in which the `srft` stations are located.

BMA forecasts for the grid locations can be obtained with `quantileForecast`, as illustrated below for the 10%, 50% (median) and 90% quantiles:

```
data(srftGrid)

memberLabels <- c("CMCG","ETA","GASP","GFS","JMA","NGPS","TCWB","UKMO")
```

```
srftGridData <- ensembleData(forecasts = srftGrid[,memberLabels],
      latitude = srftGrid[,"latitude"], longitude = srftGrid[,"longitude"],
      date = "2004013100", forecastHour = 48, initializationTime = "00")

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

The probability of freezing at grid locations can also be estimated using `cdf`, which evaluates the cumulative distribution function for the forecast model(s) corresponding to the specified date(s).

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

In datasets `srft` and `srftGrid`, temperature is recorded in degrees Kelvin (K), for which the freezing temperature corresponds to 273.15 K. The results can be displayed using the `plotProbcast` function, as shown below. Loading the `fields` (Furrer at al. 2001) and `maps` (Brownrigg and Minka 2003) libraries enables display of the country and state outlines, as well as a legend. A blue scale is chosen to display the probability of freezing, with darker shades representing higher probabilities.

```
library(fields)
library(maps)

plotProbcast( gridForc[,"0.5"], lon=srftGridData$lon,
              lat=srftGridData$lat, type="image",
              col=rev(rainbow(100,start=0,end=0.85)))
title("Median Forecast for Temperature", cex = 0.5)

bluescale <- function(n)
             hsv(4/6, s = seq(from = 1/8, to = 1, length = n), v = 1)

plotProbcast( probFreeze, lon=srftGridData$lon, lat=srftGridData$lat,
              type="image", col=bluescale(100))
title("Probability of Freezing", cex = 0.5)
```

The resulting image plots are shown in Figure 2. The plots are made by binning values onto a plotting grid. The default (shown here) is to use binning rather than interpolation to determine these values.

3.2 Precipitation Example.

In this example, we make use of the `prcpFit` and `prcpGrid` datasets included in the `ensembleBMA` package. The `prcpFit` dataset consists of the default BMA modeling parameters for the daily 48 hour forecasts of 24 hour accumulated precipitation (quantized to hundredths of an inch) over the US Pacific Northwest region from December 12, 2002

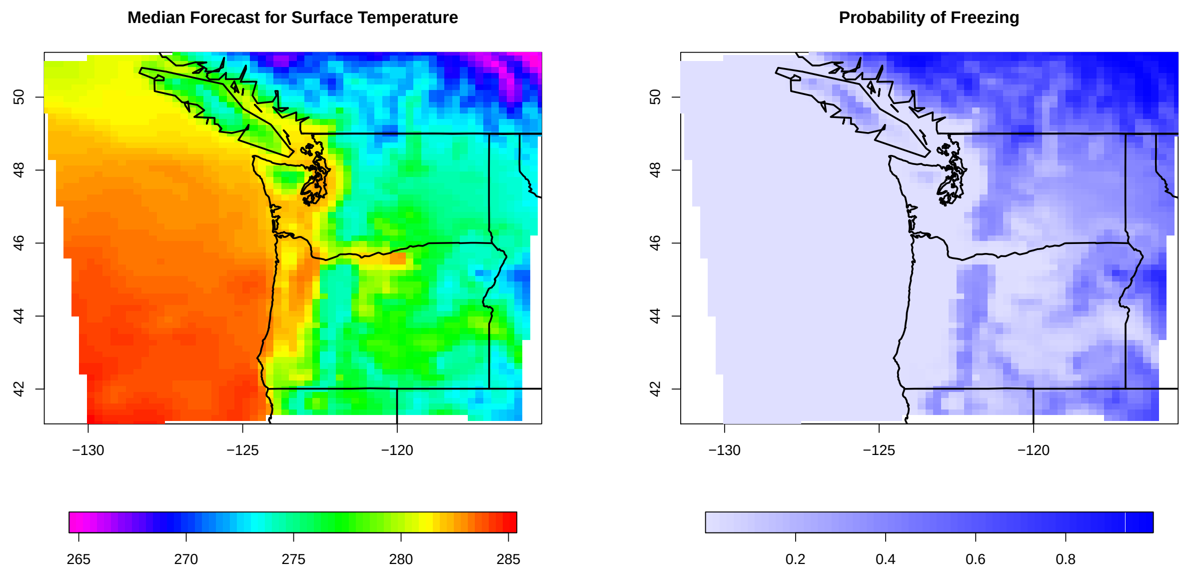

Figure 2: Image plots of the median BMA forecast of surface temperature and probability of freezing for January 31, 2004 from the `srftGrid` dataset. The plots were obtained by binning the forecasts at the grid locations onto a plotting grid. The `fields` and `maps` libraries are used to allow addition of the legend and map outline to the plot.

through March 31, 2005 used in Sloughter et al. (2007). The fitted models are mixtures of gamma distributions with a point mass at zero to the cube root transformation of the ensemble forecasts and observed data. In this case a training period of 30 days was used. The data used to obtain `prcpFit` is not included in the `ensembleBMA` package on account of its size. The `prcpGrid` dataset consists of a grid of precipitation forecasts in the region of the observations used for `prcpFit` initialized on January 13, 2003 and valid for January 15, 2003.

```
data(prcpGrid)

prcpGridData <- ensembleData(forecasts = prcpGrid[,1:9],
    latitude = prcpGrid[,"latitude"], longitude = prcpGrid[,"longitude"],
    date = "20030115", forecastHour = 48, initializationTime = "00")
```

The median and upper bound (90th percentile) forecasts can be obtained and plotted as follows:

```
data(prcpFit)

gridForc <- quantileForecast( prcpFit, prcpGridData,
                              date = "20030115", q = c(0.5, 0.9))
max(gridForc) # used to determine zlim in plotting
# 277.6505

library(fields)
```

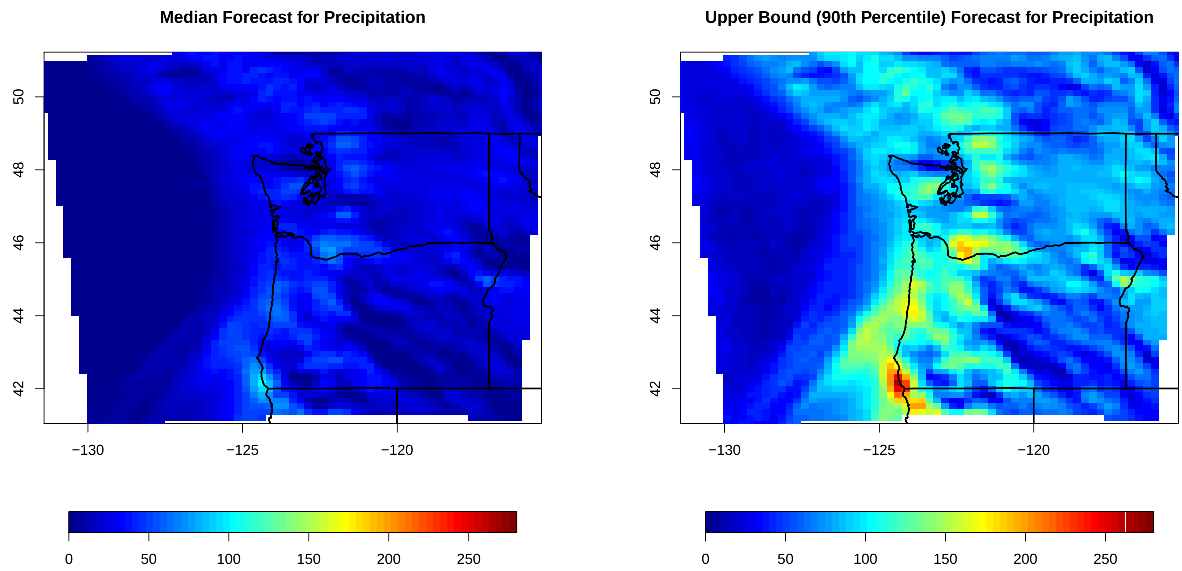

Figure 3: Image plots of the median and upper bound (90th percentile) BMA forecast of precipitation (measured in hundredths of an inch) for January 13, 2003 from the `prcpGrid` dataset. The plots were obtained by binning the forecasts at the grid locations onto a plotting grid. The `fields` and `maps` libraries are used to allow addition of the legend and map outline to the plot.

```
library(maps)

plotProbcast( gridForc[,"0.5"], type = "image", zlim = c(0,280),
              lon = prcpGridData$lon, lat = prcpGridData$lat)
title("Median Forecast for Precipitation", cex = 0.5)

plotProbcast( gridForc[,"0.9"], type = "image", zlim = c(0,280),
              lon = prcpGridData$lon, lat = prcpGridData$lat)
title("Upper Bound Forecast for Precipitation", cex = 0.5)
```

The corresponding plots are shown in Figure 3. The probability of precipitation and probability of precipitation above .25 inches can be obtained and plotted as follows. This gives an example of grayscale plotting of the data:

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

library(fields)

grayscale <- function(n) gray((0:n)/n)

range(probPrecip) # used to determine zlim in plots
# 0.02460958 0.99553477
```

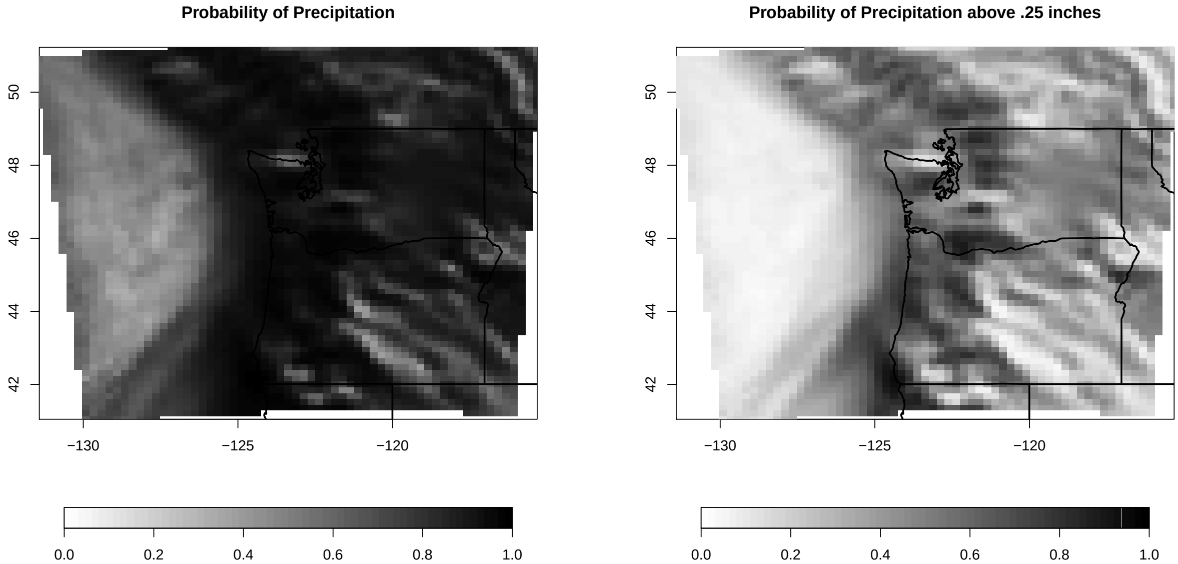

Figure 4: Grayscale image plots showing probability of precipitation for January 15, 2003 from the `prcpGrid` dataset. The plots were obtained by binning the forecasts at the grid locations onto a plotting grid. The `fields` and `maps` libraries are used to allow addition of the legend and map outline to the plot.

```
> plotProbcast( probPrecip[,"0"],
                lon=prcpGridData$lon, lat=prcpGridData$lat,
                 type="image", col= rev(grayscale(100)), zlim = c(0,1))
> title("Probability of Precipitation", cex = 0.5)

> plotProbcast( probPrecip[,"25"],
                lon=prcpGridData$lon, lat=prcpGridData$lat,
                type="image", col=rev(grayscale(100)), zlim = c(0,1))
> title("Probability of Precipitation > .25 in", cex = 0.5)
```

The corresponding plots are shown in Figure 4.

4 Verification of BMA Forecasts

The `ensembleBMA` package also provides a number of functions for verification. These can be applied to any ensemble forecasts for which both a BMA model and observed weather conditions are available. Included are functions to compute verification rank, probability integral transform, mean absolute error, continuous ranked probability scores, and Brier scores.

Surface Temperature Example. In the previous section, we obtained a forecast of surface temperature on a grid of locations for January 31, 2004 from BMA modeling of station forecasts and observations from the `srft` data set provided in the `ensembleBMA` package.

Forecasts can be obtained at the station locations by applying `quantileForecast` to the model fit `srftFit` from the previous section to the data used to generate the model.

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

These forecasts can be plotted using `plotProbcast`. The example below shows contour plots in which the `R` core function `loess` has been used to interpolate the results at the station locations onto a grid for surface plotting.

```
use <- as.character(srftData$dates) == "2004013100"
lat <- srftData$latitude[usw]; lon <- srftData$longitude[use]
lonRange <- range(lon); latRange <- range(lat)

range(srftForc[,"0.5"]) # used to determine contour levels
# 260.1318 284.7735
% 265.1425 282.0040

color <- "brown"; mapColor <- "black"

library(fields)
library(maps)

plotProbcast( srftForc[,"0.5"], lon, lat, interpolate = TRUE, col = color,
              type = "contour", levels = seq(from=264, to=284, by=2))
title("Median Forecast")
points(lon, lat, pch = 16, cex = 0.5, col = color)  # observation locations

plotProbcast( srftData$obs[use], lon, lat, interpolate = TRUE, col = color,
              type = "contour", levels = seq(from=264, to=284, by=2))
title("Observed Surface Temperature")
points(lon, lat, pch = 16, cex = 0.5, col = color)
```

The resulting plot is shown in Figure 5. In this case interpolation was used because the station locations are too sparse for binning. It is also possible to specify image or perspective plots, as well as contour plots. If the `fields` and `maps` libraries are loaded, image plots will be enhanced as shown in the displays of the previous section.

The mean continuous ranked probability score (CRPS) and mean absolute error (MAE) (see, e.g. Gneiting and Raftery (2007)) can be obtained via functions `CRPS` and `MAE`:

```
 CRPS( srftFit, srftData)
#ensemble      BMA
#1.590730 1.380526
%1.945544 1.490725

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

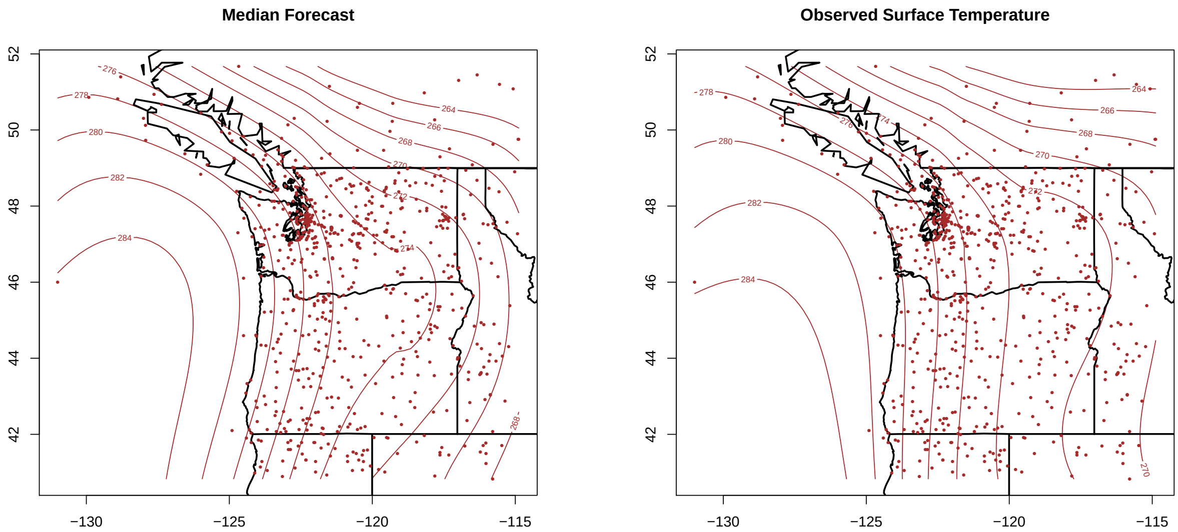

Figure 5: Contour plots of the BMA forecasts of surface temperature and verification observations at station locations for January 31, 2004 for the `srft` dataset. The plots were obtained using a `loess` fit to the forecasts and observations at the stations, interpolated on a plotting grid. The dots represent the 715 observation locations.

```
#1.929339 1.869166
%2.164045 2.042603
```

Here we are evaluating these measures for modeling at a single date; however, the CRPS and MAE would more typically be assessed over a range of dates and the corresponding models. Function `MAE` computes the mean absolute difference of the BMA median forecast [1] and the observations. The continuous ranked probability score for each observation location can be obtained using the function `crps`.

Assessing Calibration. Calibration refers to the statistical consistency between the forecast probability distribution functions and the observations (e.g. Gneiting et al. 2007). A *verification rank* histogram (e.g. Hamill 2001) can be used to assess calibration for the ensemble, while a *probability integral transform* (PIT) histogram can be used to assess calibration for the BMA forecast distributions.

The verification rank histogram plots the rank of each observation relative to the ensemble forecasts, that is, the number forecasts that are greater than the corresponding observation. It allows visual assessment of the calibration of the ensemble members. If the observation and the ensemble members come from the same distribution, then the observed and forecasted values would be exchangeable so that all possible ranks would be equally likely. We illustrate this with the surface temperature data, starting at Januray 30, 2004 — the first date for

[1] Raftery et al. (2005) use the BMA predictive mean for BMA mixtures of normals instead of the median forecast.

which we be able to postprocess forecasts with a 25 day training period using this data [2].

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

verifRankHist( ensembleForecasts(srftData[use,]),
               dataVerifObs(srftData[use,]))
```

The resulting plot is shown in Figure 6. The horizontal line shows the height that the histogram would display if the ensemble members were exchangeable. For this surface temperature data, the verification rank histogram shows a lack of calibration in the form of underdispersion for the raw ensemble.

The PIT is the value that the predictive cumulative distribution function attains at the observation and is a continuous analog of the verification rank. A function `pit` is provided in `ensembleBMA` for computing the PIT. The PIT histogram allows visual assessment of the calibration of the BMA forecasts, and is the continuous analog of 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 with corresponding data entries begin on January 30, 2004 and end on February 28, 2004):

```
# this takes time...
srftFITall <- ensembleBMA( srftData, model = "normal", trainingDays = 25)

srftPIT <- pitHist( srftFITall, srftData)
```

The resulting plot is shown in Figure 6. The horizontal line shows the height that the histogram would display if the ensemble members were exchangeable. For this surface temperature data, the PIT histogram shows signs of negative bias, which is not surpising 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.

Precipitation Example. In the previous section, we obtained a forecast of precipitation on a grid of locations for January 15, 2003 from BMA modeling of station forecasts and observations from the `prcpFit` and `prcpGrid` datasets provided in the `ensembleBMA` package. Quantile forecasts can be obtained at the station locations by applying `quantileForecast` to the model fit given the data used to generate the model. An `ensembleBMA` object called `prcpDJdata` is provided as a dataset with the package containing ensemble forecasts and verification observations for this date. We can compare the forecasts with the observed data graphically using the function `verifPlot` as follows:

```
> data(prcpDJdata)
> forc <- verifPlot( prcpFit, prcpDJdata, date = "20030113")
```

The resulting plot is shown in Figure 7.

[2]Although datasets with missing values can be handled in `ensembleBMA`, the `srft` dataset has no missing values, and ensemble forecasts are missing for some dates.

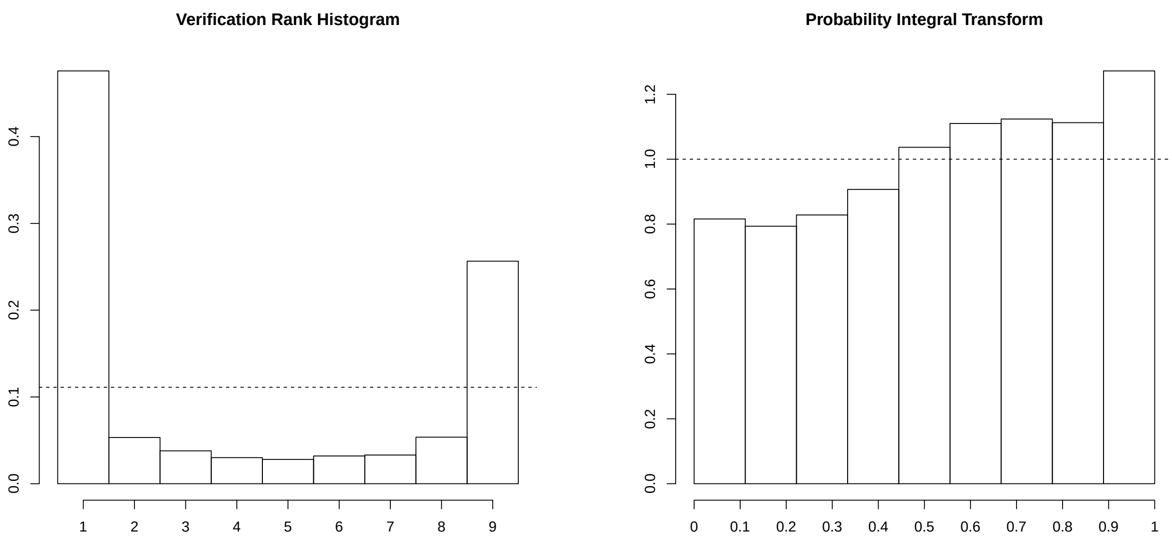

Figure 6: Verification rank histogram for the observed surface temperature relative to the ensemble for January 30, 2004, and PIT histogram for the observed surface temperature relative to the BMA forecasts for all of the `srft` data. A more uniform histogram implies better calibration, with the dotted lines indicating the histogram height corresponding to perfect calibration. The area under the histogram is equal to 1 in each case, with the heights differing because of the difference in the scale of the horizontal axes.

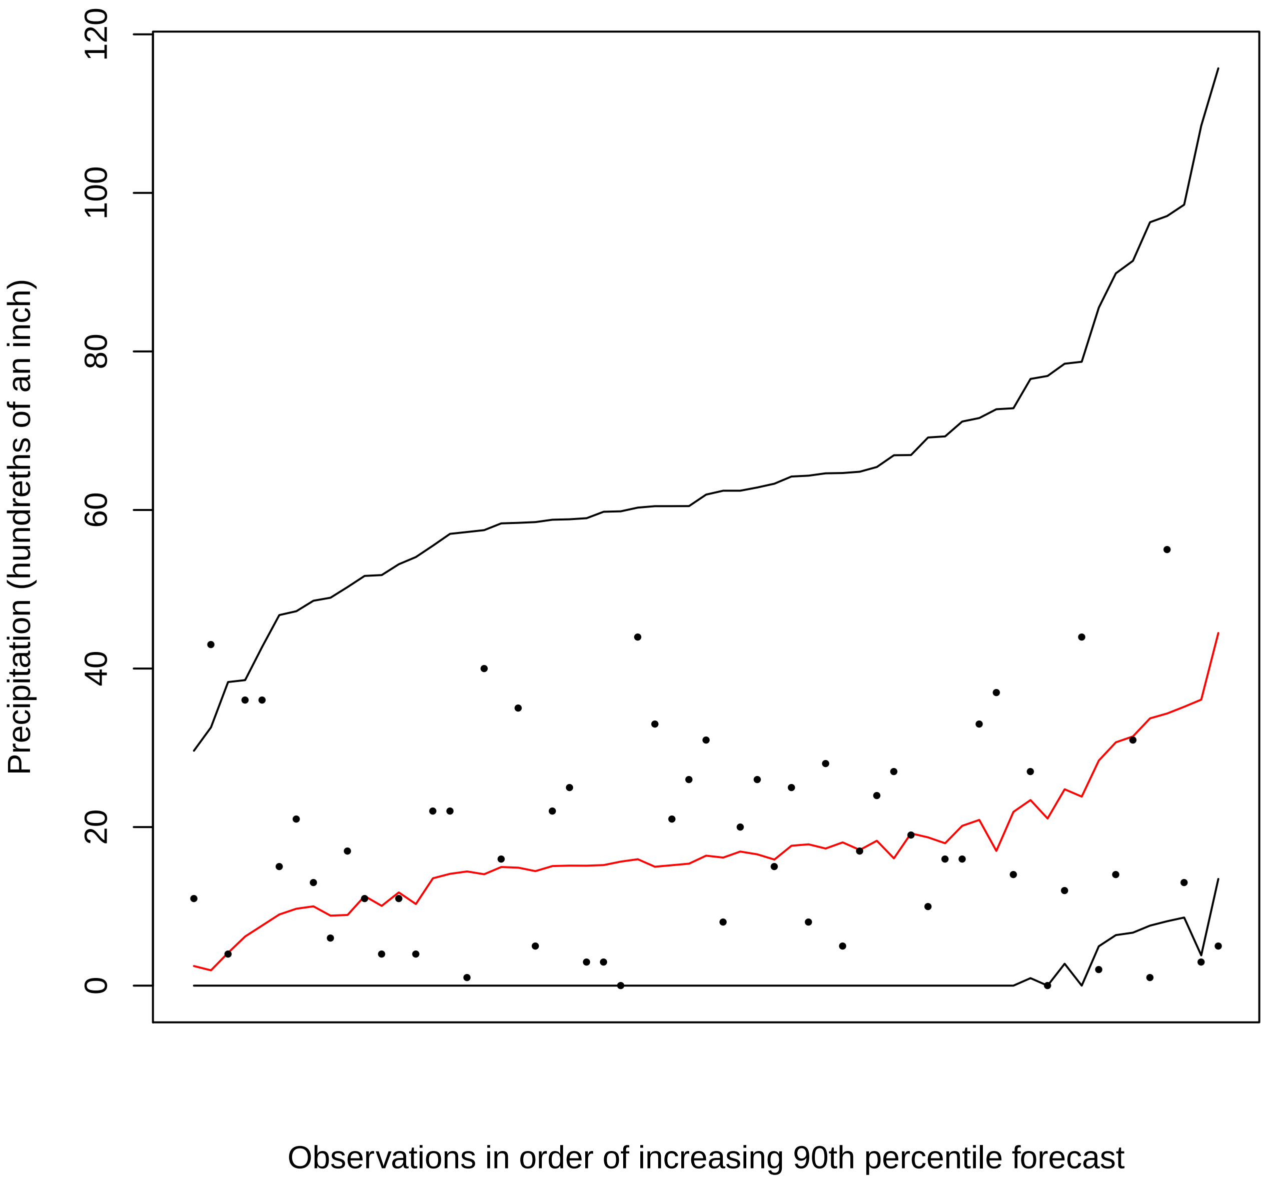

Figure 7: The lines represent the 10th (gray), 50th (red), and 90th (black) percentile BMA forecasts of precipitation for January 15, 2003 at the station locations, while the dots indicate the observed precipitation at the same locations. The horizontal axis represents the observations, in order of increasing 90th percentile forecast.

The mean continuous ranked probability score (CRPS) and mean absolute error (MAE) can be obtained via functions `CRPS` and `MAE`. Here we have done so for the entire precipitation data set available from `http://www.stat.washington.edu/MURI`. It is not included in the `ensembleBMA` package on account of its size.

```
 CRPS( prcpFit, prcpDJdata)
#ensemble      BMA
#13.78853 12.18904

  MAE( prcpFit, prcpDJdata)
#ensemble      BMA
#17.49649 16.43995
```

For BMA mixtures of gammas with a point mass at 0, `MAE` computes the mean absolute difference of the BMA median forecast and the observations (Sloughter et al. 2007). Brier scores (see, e.g. Joliffe and Stephenson, 2003) for the model fits can be obtained using `brierScore`.

```
 brierScore( prcpFit, prcpDJdata, thresh = c(0, 50, 100, 200, 300, 400))
#    thresholds climatology    ensemble    logistic         bma
#0            0 0.234100222 0.179320900 0.121370659 0.125871040
#50          50 0.121871990 0.092102740 0.076341763 0.083166151
#100        100 0.045193270 0.039034799 0.033518508 0.035444253
#200        200 0.009786516 0.010591972 0.008798700 0.008873115
#300        300 0.004498081 0.004982201 0.004409984 0.004298522
#400        400 0.002816086 0.003134359 0.002816054 0.002840496
```

Here 'climatology' refers to the empirical distribution of the verifying observations, while 'logistic' refers a logistic regression model with the cube root of the data as predictor variable, with coefficients determined from the training data. This logistic regression model is the one used for the probability of precipitation component in the forecasting model of Sloughter et al. (2007).

5 Function Summary

The main functions in the `ensembleBMA` package are:

`ensembleData`: Creates a data object with forecasts and (optionally) observations, dates and other indentifying information for the forecasts and observations.

`ensembleBMA`: Fits one or more BMA models to forecasting data given a training rule defined by the number of training days and the forecast hour.

`fitBMA`: Fits a BMA model to training data.

Given a model created by `ensembleBMA` or `fitBMA` and corresponding ensemble forecasts (need not be the same data used in modeling): `priorBMAgamma0`: Computes prior values for logistic regression coefficients for probability of zero percipitation in the `gamma0` averaging their values over a number of training periods.

Given a model created by `ensembleBMA` or `fitBMA` and corresponding ensemble forecasts (need not be the same data used in modeling):

> `quantileForecast`: Determines probabilistic forecasts from the BMA models.
>
> `cdf`: Computes the cumulative distribution function for the BMA models.
>
> `combine`: Combines compatible `ensembleBMA` models that have different dates.
>
> `crps`: Computes continuous rank probability scores for the ensemble and BMA.
>
> `CRPS`: Computes the mean continuous rank probability scores.
>
> `MAE`: Computes the mean absolute error for the ensemble and BMA median forecasts.
>
> `plot`: Plots the predictive distribution at a each location and each date.
>
> `pit`: Computes the probability integral transform of the observations given the BMA model; that is, the value of the cumulative distribution function at the observations.
>
> `pitHist`: Returns the probability intergral transform of the observations given the BMA model and plots its histogram.
>
> `verifPlot`: Returns the median, 10th and 90th percentile forecasts, and plots them along with the observations in order of increasing 90th percentile forecast.
>
> `verifRankHist`: Returns the verification rank of the verifying observations relative to the corresponding ensemble forecasts, and plots its histogram.
>
> `brierScore`: Computes Brier Scores for the ensemble and for BMA.

Other functions in the `ensembleBMA` package:

`modelParameters`: Extracts model parameters from `ensembleBMA` or `fitBMA` output.

`controlBMAgamma,controlBMAgamma0,controlBMAnormal` Functions for setting values controling model fitting in `ensembleBMA` and `fitBMA`.

`plotProbcast`: Function for plotting forecasts.

`trainingData`: Extracts training data from an `ensembleData` object corresponding to a given number of training days for a single specified date.

`dateCheck`: Checks that dates correspond to `YYYYMMDDHH` or `YYYYMMDD` format.

`ymdhTOjul`: Converts `YYYYMMDDHH` or `YYYYMMDD` formatted dates to Julian dates.

`julTOymdh`: Converts Julian dates to `YYYYMMDDHH` or `YYYYMMDD` formatted dates.

References

[1] 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.

[2] V. J. Berrocal, A. E. Raftery, T. Gneiting, and R. C. Steed. Probabilistic weather forecasting for winter road maintenance. *Journal of the American Statistical Association*, 105:522–537, 2010.

[3] R. Brownrigg and T. P. Minka. maps: Draw geographical maps, 2003. (R package, latest revision 2010, original S code by R. A. Becker and A. R. Wilks).

[4] F. A. Eckel and C. F. Mass. Effective mesoscale, short-range ensemble forecasting. *Weather and Forecasting*, 20:328–350, 2005.

[5] 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.

[6] R. Furrer, D. Nychka, and S. Sain. fields: Tools for spatial data, 2001. (R package, latest revision 2009).

[7] 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.

[8] T. Gneiting and A. E. Raftery. Weather forecasting with ensemble methods. *Science*, 310:248–249, 2005.

[9] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.

[10] 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.

[11] T. M. Hamill. Interpretation of rank histograms for verifying ensemble forecasts. *Monthly Weather Review*, 129:550–560, 2001.

[12] K. Hornik. chron: Chronological objects which can handle dates and times, 1999. (R package, latest revision 2009, S original by D. James).

[13] I. T. Jolliffe and D. B. Stephenson, editors. *Forecast Verification: A Practitioner's Guide in Atmospheric Science*. Wiley, 2003.

[14] 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.

[15] 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, 2009.

[16] 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.

[17] 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.