

Package ‘opera’

December 6, 2021

Type Package

Title Online Prediction by Expert Aggregation

Version 1.2.0

Author Pierre Gaillard [cre, aut],
Yannig Goude [aut],
Laurent Plagne [ctb],
Thibaut Dubois [ctb],
Benoit Thieurmél [ctb]

Maintainer Pierre Gaillard <pierre@gaillard.me>

Copyright EDF R&D 2012-2015

Description Misc methods to form online predictions, for regression-oriented time-series, by combining a finite set of forecasts provided by the user. See Cesa-Bianchi and Lugosi (2006) <[doi:10.1017/CBO9780511546921](https://doi.org/10.1017/CBO9780511546921)> for an overview.

License LGPL

URL <http://pierre.gaillard.me/opera.html>

BugReports <https://github.com/dralliag/opera/issues>

Depends R (>= 3.5.0)

LinkingTo Rcpp, RcppEigen

Imports Rcpp, htmltools, rAmCharts, htmlwidgets, pipeR, alabama, methods, Rdpack

Suggests quantreg, quadprog, RColorBrewer, testthat, splines, caret, mgcv, survival, knitr, gbm, rmarkdown, magrittr

RdMacros Rdpack

LazyData true

VignetteBuilder knitr

RoxygenNote 7.1.2

Encoding UTF-8

NeedsCompilation yes

Repository CRAN

Date/Publication 2021-12-06 18:00:02 UTC

R topics documented:

opera-package	2
check_loss	5
check_matrix	6
electric_load	7
FTRL	7
loss	9
mixture	10
oracle	17
plot.mixture	19
plot.oracle	20
plot_ridge_weights	21
plt_oracle_convex	22
predict.mixture	22
seriesToBlock	23

Index	25
--------------	-----------

opera-package	<i>Online Prediction by ExpeRt Aggregation</i>
---------------	--

Description

The package `opera` performs, for regression-oriented time-series, predictions by combining a finite set of forecasts provided by the user. More formally, it considers a sequence of observations Y (such as electricity consumption, or any bounded time series) to be predicted step by step. At each time instance t , a finite set of experts (basically some based forecasters) provide predictions x of the next observation in y . This package proposes several adaptive and robust methods to combine the expert forecasts based on their past performance.

Author(s)

Pierre Gaillard <pierre@gaillard.me>

References

Prediction, Learning, and Games. N. Cesa-Bianchi and G. Lugosi.

Forecasting the electricity consumption by aggregating specialized experts; a review of sequential aggregation of specialized experts, with an application to Slovakian and French country-wide one-day-ahead (half-)hourly predictions, *Machine Learning*, in press, 2012. Marie Devaine, Pierre Gaillard, Yannig Goude, and Gilles Stoltz

Contributions to online robust aggregation: work on the approximation error and on probabilistic forecasting. Pierre Gaillard. PhD Thesis, University Paris-Sud, 2015.

Examples

```

library('opera') # load the package
set.seed(1)

# Example: find the best one week ahead forecasting strategy (weekly data)
# packages
library(mgcv)

# import data
data(electric_load)
idx_data_test <- 620:nrow(electric_load)
data_train <- electric_load[-idx_data_test, ]
data_test <- electric_load[idx_data_test, ]

# Build the expert forecasts
# #####

# 1) A generalized additive model
gam.fit <- gam(Load ~ s(IPI) + s(Temp) + s(Time, k=3) +
              s(Load1) + as.factor(NumWeek), data = data_train)
gam.forecast <- predict(gam.fit, newdata = data_test)

# 2) An online autoregressive model on the residuals of a medium term model

# Medium term model to remove trend and seasonality (using generalized additive model)
detrrend.fit <- gam(Load ~ s(Time,k=3) + s(NumWeek) + s(Temp) + s(IPI), data = data_train)
electric_load$Trend <- c(predict(detrrend.fit), predict(detrrend.fit,newdata = data_test))
electric_load$Load.detrrend <- electric_load$Load - electric_load$Trend

# Residual analysis
ar.forecast <- numeric(length(idx_data_test))
for (i in seq(idx_data_test)) {
  ar.fit <- ar(electric_load$Load.detrrend[1:(idx_data_test[i] - 1)])
  ar.forecast[i] <- as.numeric(predict(ar.fit)$pred) + electric_load$Trend[idx_data_test[i]]
}

# Aggregation of experts
#####

X <- cbind(gam.forecast, ar.forecast)
colnames(X) <- c('gam', 'ar')
Y <- data_test$Load

matplot(cbind(Y, X), type = 'l', col = 1:6, ylab = 'Weekly load', xlab = 'Week')

# How good are the expert? Look at the oracles
oracle.convex <- oracle(Y = Y, experts = X, loss.type = 'square', model = 'convex')

if(interactive()){
  plot(oracle.convex)
}

```

```

oracle.convex

# Is a single expert the best over time ? Are there breaks ?
oracle.shift <- oracle(Y = Y, experts = X, loss.type = 'percentage', model = 'shifting')
if(interactive()){
  plot(oracle.shift)
}
oracle.shift

# Online aggregation of the experts with BOA
#####

# Initialize the aggregation rule
m0.BOA <- mixture(model = 'BOA', loss.type = 'square')

# Perform online prediction using BOA There are 3 equivalent possibilities 1)
# start with an empty model and update the model sequentially
m1.BOA <- m0.BOA
for (i in 1:length(Y)) {
  m1.BOA <- predict(m1.BOA, newexperts = X[i, ], newY = Y[i], quiet = TRUE)
}

# 2) perform online prediction directly from the empty model
m2.BOA <- predict(m0.BOA, newexpert = X, newY = Y, online = TRUE, quiet = TRUE)

# 3) perform the online aggregation directly
m3.BOA <- mixture(Y = Y, experts = X, model = 'BOA', loss.type = 'square', quiet = TRUE)

# These predictions are equivalent:
identical(m1.BOA, m2.BOA) # TRUE
identical(m1.BOA, m3.BOA) # TRUE

# Display the results
summary(m3.BOA)
if(interactive()){
  plot(m1.BOA)
}

# Using d-dimensional time-series
#####

# Consider the above exemple of electricity consumption
# to be predicted every four weeks
YBlock <- seriesToBlock(X = Y, d = 4)
XBlock <- seriesToBlock(X = X, d = 4)

# The four-week-by-four-week predictions can then be obtained
# by directly using the `mixture` function as we did earlier.

MLpolBlock <- mixture(Y = YBlock, experts = XBlock, model = "MLpol", loss.type = "square",
  quiet = TRUE)

```

```

# The predictions can finally be transformed back to a
# regular one dimensional time-series by using the function `blockToSeries`.

prediction <- blockToSeries(MLpolBlock$prediction)

#### Using the `online = FALSE` option

# Equivalent solution is to use the `online = FALSE` option in the predict function.
# The latter ensures that the model coefficients are not
# updated between the next four weeks to forecast.
MLpolBlock <- mixture(model = "BOA", loss.type = "square")
d = 4
n <- length(Y)/d
for (i in 0:(n-1)) {
  idx <- 4*i + 1:4 # next four weeks to be predicted
  MLpolBlock <- predict(MLpolBlock, newexperts = X[idx, ], newY = Y[idx], online = FALSE,
                        quiet = TRUE)
}

print(head(MLpolBlock$weights))

```

check_loss

Function to check validity of provided loss function

Description

Function to check validity of provided loss function

Usage

```

check_loss(
  loss.type,
  loss.gradient,
  Y = NULL,
  model = NULL,
  use_cpp = getOption("opera_use_cpp", default = FALSE)
)

```

Arguments

loss.type character, list or function.

character Name of the loss to be applied ('square', 'absolute', 'percentage', or 'pinball');

list When using pinball loss: list with field name equal to 'pinball' and field tau equal to the required quantile in [0,1];

function A custom loss as a function of two parameters.

loss.gradient	boolean, function. boolean If TRUE, the aggregation rule will not be directly applied to the loss function at hand, but to a gradient version of it. The aggregation rule is then similar to gradient descent aggregation rule. function If loss.type is a function, the derivative should be provided to be used (it is not automatically computed).
Y	numeric (NULL). (Optional) Target values (to perform some checks).
model	character (NULL). (Optional) Model used (to perform some checks).
use_cpp	boolean. Whether or not to use cpp function to increase perf.

Value

loss.type

check_matrix *Function to check and modify the input class and type*

Description

Function to check and modify the input class and type

Usage

```
check_matrix(mat, name)
```

Arguments

mat	data.frame, data.table, tibble. Object to be cast to matrix.
name	character. Name of the object to be cast.

Value

a 3d array if a 3d array is provided, else a matrix.

electric_load	<i>Electricity forecasting data set</i>
---------------	---

Description

Electricity forecasting data set provided by EDF R&D. It contains weekly measurements of the total electricity consumption in France from 1996 to 2009, together with several covariates, including temperature, industrial production indices (source: INSEE) and calendar information.

Usage

```
data(electric_load)
```

Format

An object of class `data.frame` with 731 rows and 11 columns.

Examples

```
data(electric_load)
# a few graphs to display the data
attach(electric_load)
plot(Load, type = 'l')
plot(Temp, Load, pch = 16, cex = 0.5)
plot(NumWeek, Load, pch = 16, cex = 0.5)
plot(Load, Load1, pch = 16, cex = 0.5)
acf(Load, lag.max = 20)
detach(electric_load)
```

FTRL	<i>Implementation of FTRL (Follow The Regularized Leader)</i>
------	---

Description

FTRL (Shalev-Shwartz and Singer 2007) and Chap. 5 of (Hazan 2019) is the online counterpart of empirical risk minimization. It is a family of aggregation rules (including OGD) that uses at any time the empirical risk minimizer so far with an additional regularization. The online optimization can be performed on any bounded convex set that can be expressed with equality or inequality constraints. Note that this method is still under development and a beta version.

Usage

```

FTRL(
  y,
  experts,
  eta = NULL,
  fun_reg = NULL,
  fun_reg_grad = NULL,
  constr_eq = NULL,
  constr_eq_jac = NULL,
  constr_ineq = NULL,
  constr_ineq_jac = NULL,
  loss.type = list(name = "square"),
  loss.gradient = TRUE,
  w0 = NULL,
  max_iter = 50,
  obj_tol = 0.01,
  training = NULL,
  default = FALSE,
  quiet = TRUE
)

```

Arguments

<code>y</code>	vector. Real observations.
<code>experts</code>	matrix. Matrix of experts previsions.
<code>eta</code>	numeric. Regularization parameter.
<code>fun_reg</code>	function (NULL). Regularization function to be applied during the optimization.
<code>fun_reg_grad</code>	function (NULL). Gradient of the regularization function (to speed up the computations).
<code>constr_eq</code>	function (NULL). Constraints (equalities) to be applied during the optimization.
<code>constr_eq_jac</code>	function (NULL). Jacobian of the equality constraints (to speed up the computations).
<code>constr_ineq</code>	function (NULL). Constraints (inequalities) to be applied during the optimization (... > 0).
<code>constr_ineq_jac</code>	function (NULL). Jacobian of the inequality constraints (to speed up the computations).
<code>loss.type</code>	character, list or function ("square"). character Name of the loss to be applied ('square', 'absolute', 'percentage', or 'pinball'); list List with field name equal to the loss name. If using pinball loss, field tau equal to the required quantile in [0,1]; function A custom loss as a function of two parameters (prediction, label).

loss.gradient	boolean, function (TRUE). boolean If TRUE, the aggregation rule will not be directly applied to the loss function at hand, but to a gradient version of it. The aggregation rule is then similar to gradient descent aggregation rule. function If loss.type is a function, the derivative of the loss in its first component should be provided to be used (it is not automatically computed).
w0	numeric (NULL). Vector of initialization for the weights.
max_iter	integer (50). Maximum number of iterations of the optimization algorithm per round.
obj_tol	numeric (1e-2). Tolerance over objective function between two iterations of the optimization.
training	list (NULL). List of previous parameters.
default	boolean (FALSE). Whether or not to use default parameters for fun_reg, constr_eq, constr_ineq and their grad/jac, which values are ALL ignored when TRUE.
quiet	boolean (FALSE). Whether or not to display progress bars.

Value

object of class mixture.

References

Hazan E (2019). “Introduction to online convex optimization.” *arXiv preprint arXiv:1909.05207*.

Shalev-Shwartz S, Singer Y (2007). “A primal-dual perspective of online learning algorithms.” *Machine Learning*, **69**(2), 115–142.

loss *Errors suffered by a sequence of predictions*

Description

The function loss computes the sequence of instantaneous losses suffered by the predictions in x to predict the observation in y .

Usage

```
loss(
  x,
  y,
  pred = NULL,
  loss.type = list(name = "square"),
  loss.gradient = FALSE
)
```

Arguments

- `x` numeric. A vector of length T containing the sequence of prediction to be evaluated.
- `y` numeric. A vector of length T that contains the observations to be predicted.
- `pred` numeric. A vector of length T containing the sequence of real values.
- `loss.type` character, list or function ("square").
- character Name of the loss to be applied ('square', 'absolute', 'percentage', or 'pinball');
 - list List with field name equal to the loss name. If using pinball loss, field `tau` equal to the required quantile in $[0,1]$;
 - function A custom loss as a function of two parameters.
- `loss.gradient` boolean, function (TRUE).
- boolean If TRUE, the aggregation rule will not be directly applied to the loss function at hand, but to a gradient version of it. The aggregation rule is then similar to gradient descent aggregation rule.
 - function If `loss.type` is a function, the derivative should be provided to be used (it is not automatically computed).

Value

A vector of length T containing the sequence of instantaneous losses suffered by the expert predictions (`x`) or the gradient computed on the aggregated predictions (`pred`).

Author(s)

Pierre Gaillard <pierre@gaillard.me>

mixture

Compute an aggregation rule

Description

The function `mixture` builds an aggregation rule chosen by the user. It can then be used to predict new observations Y sequentially. If observations Y and expert advice `experts` are provided, `mixture` is trained by predicting the observations in Y sequentially with the help of the expert advice in `experts`. At each time instance $t = 1, 2, \dots, T$, the mixture forms a prediction of $Y[t,]$ by assigning a weight to each expert and by combining the expert advice.

Usage

```
mixture(
  Y = NULL,
  experts = NULL,
  model = "MLpol",
```

```

    loss.type = "square",
    loss.gradient = TRUE,
    coefficients = "Uniform",
    awake = NULL,
    parameters = list(),
    use_cpp = getOption("opera_use_cpp", default = FALSE),
    quiet = TRUE
)

## S3 method for class 'mixture'
print(x, ...)

## S3 method for class 'mixture'
summary(object, ...)

```

Arguments

Y	A matrix with T rows and d columns. Each row $Y[t,]$ contains a d-dimensional observation to be predicted sequentially.
experts	An array of dimension $c(T, d, K)$, where T is the length of the data-set, d the dimension of the observations, and K is the number of experts. It contains the expert forecasts. Each vector $experts[t, k]$ corresponds to the d-dimensional prediction of $Y[t,]$ proposed by expert k at time $t = 1, \dots, T$. In the case of real prediction (i.e., $d = 1$), <i>experts</i> is a matrix with T rows and K columns.
model	A character string specifying the aggregation rule to use. Currently available aggregation rules are: <ul style="list-style-type: none"> 'EWA' Exponentially weighted average aggregation rules (Cesa-Bianchi and Lugosi 2006). A positive learning rate eta can be chosen by the user. The bigger it is the faster the aggregation rule will learn from observations and experts performances. However, too high values lead to unstable weight vectors and thus unstable predictions. If it is not specified, the learning rate is calibrated online. A finite grid of potential learning rates to be optimized online can be specified with grid.eta. 'FS' Fixed-share aggregation rule (Cesa-Bianchi and Lugosi 2006). As for ewa, a learning rate eta can be chosen by the user or calibrated online. The main difference with ewa aggregation rule rely in the mixing rate alpha $\in [0, 1]$ which considers at each instance a small probability alpha to have a rupture in the sequence and that the best expert may change. Fixed-share aggregation rule can thus compete with the best sequence of experts that can change a few times (see oracle), while ewa can only compete with the best fixed expert. The mixing rate alpha is either chosen by the user either calibrated online. Finite grids of learning rates and mixing rates to be optimized can be specified with parameters grid.eta and grid.alpha. 'Ridge' Online Ridge regression (Cesa-Bianchi and Lugosi 2006). It minimizes at each instance a penalized criterion. It forms at each instance linear combination of the experts' forecasts and can assign negative weights that not necessarily sum to one. It is useful if the experts are biased or correlated. It cannot be used with specialized experts. A positive regularization

coefficient **lambda** can either be chosen by the user or calibrated online. A finite grid of coefficient to be optimized can be specified with a parameter **grid.lambda**.

'**MLpol**', '**MLewa**', '**MLprod**' Aggregation rules with multiple learning rates that are theoretically calibrated (Gaillard et al. 2014).

'**BOA**' Bernstein online Aggregation (Wintenberger 2017). The learning rates are automatically calibrated.

'**OGD**' Online Gradient descent (Zinkevich 2003). See also (Hazan 2019). The optimization is performed with a time-varying learning rate. At time step $t \geq 1$, the learning rate is chosen to be $t^{-\alpha}$, where α is provided by alpha in the parameters argument. The algorithm may or not perform a projection step into the simplex space (non-negative weights that sum to one) according to the value of the parameter 'simplex' provided by the user.

'**FTRL**' Follow The Regularized Leader (Shalev-Shwartz and Singer 2007). Note that here, the linearized version of FTRL is implemented (see Chap. 5 of (Hazan 2019)). **FTRL** is the online counterpart of empirical risk minimization. It is a family of aggregation rules (including OGD) that uses at any time the empirical risk minimizer so far with an additional regularization. The online optimization can be performed on any bounded convex set that can be expressed with equality or inequality constraints. Note that this method is still under development and a beta version.

The user must provide (in the **parameters**'s list):

- 'eta' The learning rate.
- 'fun_reg' The regularization function to be applied on the weights. See [auglag](#): fn.
- 'constr_eq' The equality constraints (e.g. $\sum(w) = 1$). See [auglag](#): heq.
- 'constr_ineq' The inequality constraints (e.g. $w > 0$). See [auglag](#): hin.
- 'fun_reg_grad' (optional) The gradient of the regularization function. See [auglag](#): gr.
- 'constr_eq_jac' (optional) The Jacobian of the equality constraints. See [auglag](#): heq.jac
- 'constr_ineq_jac' (optional) The Jacobian of the inequality constraints. See [auglag](#): hin.jac

or set **default** to **TRUE**. In the latter, **FTRL** is performed with Kullback regularization ($\text{fun_reg}(x) = \sum(x \log(x/w_0))$) on the simplex ($\text{constr_eq}(w) = \sum(w) - 1$ and $\text{constr_ineq}(w) = w$). Parameters **w0** (weight initialization), and **max_iter** can also be provided.

loss.type

character, list, or function ("square").

character Name of the loss to be applied ('square', 'absolute', 'percentage', or 'pinball');

list List with field name equal to the loss name. If using pinball loss, field tau equal to the required quantile in [0,1];

function A custom loss as a function of two parameters (prediction, observation). For example, $f(x,y) = \text{abs}(x-y)/y$ for the Mean absolute percentage error or $f(x,y) = (x-y)^2$ for the squared loss.

loss.gradient	<p>boolean,function (TRUE).</p> <p>boolean If TRUE, the aggregation rule will not be directly applied to the loss function at hand, but to a gradient version of it. The aggregation rule is then similar to gradient descent aggregation rule.</p> <p>function Can be provided if loss.type is a function. It should then be a sub-derivative of the loss in its first component (i.e., in the prediction). For instance, $g(x) = (x-y)$ for the squared loss.</p>
coefficients	A probability vector of length K containing the prior weights of the experts (not possible for 'MLpol'). The weights must be non-negative and sum to 1.
awake	A matrix specifying the activation coefficients of the experts. Its entries lie in $[0, 1]$. Possible if some experts are specialists and do not always form and suggest prediction. If the expert number k at instance t does not form any prediction of observation Y_t , we can put $awake[t, k]=0$ so that the mixture does not consider expert k in the mixture to predict Y_t .
parameters	<p>A list that contains optional parameters for the aggregation rule. If no parameters are provided, the aggregation rule is fully calibrated online. Possible parameters are:</p> <p>eta A positive number defining the learning rate. Possible if model is either 'EWA' or 'FS'</p> <p>grid.eta A vector of positive numbers defining potential learning rates for 'EWA' of 'FS'. The learning rate is then calibrated by sequentially optimizing the parameter in the grid. The grid may be extended online if needed by the aggregation rule.</p> <p>gamma A positive number defining the exponential step of extension of grid.eta when it is needed. The default value is 2.</p> <p>alpha A number in $[0, 1]$. If the model is 'FS', it defines the mixing rate. If the model is 'OGD', it defines the order of the learning rate: $\eta_t = t^{-\alpha}$.</p> <p>grid.alpha A vector of numbers in $[0, 1]$ defining potential mixing rates for 'FS' to be optimized online. The grid is fixed over time. The default value is $[0.0001, 0.001, 0.01, 0.1]$.</p> <p>lambda A positive number defining the smoothing parameter of 'Ridge' aggregation rule.</p> <p>grid.lambda Similar to grid.eta for the parameter lambda.</p> <p>simplex A boolean that specifies if 'OGD' does a project on the simplex. In other words, if TRUE (default) the online gradient descent will be under the constraint that the weights sum to 1 and are non-negative. If FALSE, 'OGD' performs an online gradient descent on K dimensional real space. without any projection step.</p> <p>averaged A boolean (default is FALSE). If TRUE the coefficients and the weights returned (and used to form the predictions) are averaged over the past. It leads to more stability on the time evolution of the weights but needs more regularity assumption on the underlying process generating the data (i.i.d. for instance).</p>
use_cpp	boolean. Whether or not to use cpp optimization to fasten the computations. This option is not yet compatible with the use of custom loss function. Note that

cpp implementation corresponds to an earlier version of the code and may be outdated. Use `options(opera_use_cpp = TRUE)` to change the default value.

quiet	boolean. Whether or not to display progress bars.
x	An object of class mixture
...	Additional parameters
object	An object of class mixture

Value

An object of class mixture that can be used to perform new predictions. It contains the parameters `model`, `loss.type`, `loss.gradient`, `experts`, `Y`, `awake`, and the fields

coefficients	A vector of coefficients assigned to each expert to perform the next prediction.
weights	A matrix of dimension $c(T, K)$, with T the number of instances to be predicted and K the number of experts. Each row contains the convex combination to form the predictions
prediction	A matrix with T rows and d columns that contains the predictions outputted by the aggregation rule.
loss	The average loss (as stated by parameter <code>loss.type</code>) suffered by the aggregation rule.
parameters	The learning parameters chosen by the aggregation rule or by the user.
training	A list that contains useful temporary information of the aggregation rule to be updated and to perform predictions.

Author(s)

Pierre Gaillard <pierre@gaillard.me> Yannig Goude <yannig.goude@edf.fr>

References

- Cesa-Bianchi N, Lugosi G (2006). *Prediction, learning, and games*. Cambridge university press.
- Gaillard P, Stoltz G, van Erven T (2014). "A Second-order Bound with Excess Losses." In *Proceedings of COLT'14*, volume 35, 176–196.
- Hazan E (2019). "Introduction to online convex optimization." *arXiv preprint arXiv:1909.05207*.
- Shalev-Shwartz S, Singer Y (2007). "A primal-dual perspective of online learning algorithms." *Machine Learning*, **69**(2), 115–142.
- Wintenberger O (2017). "Optimal learning with Bernstein online aggregation." *Machine Learning*, **106**(1), 119–141.
- Zinkevich M (2003). "Online convex programming and generalized infinitesimal gradient ascent." In *Proceedings of the 20th international conference on machine learning (icml-03)*, 928–936.

See Also

See [opera-package](#) and [opera-vignette](#) for a brief example about how to use the package.

Examples

```

library('opera') # load the package
set.seed(1)

# Example: find the best one week ahead forecasting strategy (weekly data)
# packages
library(mgcv)

# import data
data(electric_load)
idx_data_test <- 620:nrow(electric_load)
data_train <- electric_load[-idx_data_test, ]
data_test <- electric_load[idx_data_test, ]

# Build the expert forecasts
# #####

# 1) A generalized additive model
gam.fit <- gam(Load ~ s(IPI) + s(Temp) + s(Time, k=3) +
              s(Load1) + as.factor(NumWeek), data = data_train)
gam.forecast <- predict(gam.fit, newdata = data_test)

# 2) An online autoregressive model on the residuals of a medium term model

# Medium term model to remove trend and seasonality (using generalized additive model)
detrrend.fit <- gam(Load ~ s(Time,k=3) + s(NumWeek) + s(Temp) + s(IPI), data = data_train)
electric_load$Trend <- c(predict(detrrend.fit), predict(detrrend.fit,newdata = data_test))
electric_load$Load.detrrend <- electric_load$Load - electric_load$Trend

# Residual analysis
ar.forecast <- numeric(length(idx_data_test))
for (i in seq(idx_data_test)) {
  ar.fit <- ar(electric_load$Load.detrrend[1:(idx_data_test[i] - 1)])
  ar.forecast[i] <- as.numeric(predict(ar.fit)$pred) + electric_load$Trend[idx_data_test[i]]
}

# Aggregation of experts
#####

X <- cbind(gam.forecast, ar.forecast)
colnames(X) <- c('gam', 'ar')
Y <- data_test$Load

matplot(cbind(Y, X), type = 'l', col = 1:6, ylab = 'Weekly load', xlab = 'Week')

# How good are the expert? Look at the oracles
oracle.convex <- oracle(Y = Y, experts = X, loss.type = 'square', model = 'convex')

```

```

if(interactive()){
  plot(oracle.convex)
}

oracle.convex

# Is a single expert the best over time ? Are there breaks ?
oracle.shift <- oracle(Y = Y, experts = X, loss.type = 'percentage', model = 'shifting')
if(interactive()){
  plot(oracle.shift)
}
oracle.shift

# Online aggregation of the experts with BOA
#####

# Initialize the aggregation rule
m0.BOA <- mixture(model = 'BOA', loss.type = 'square')

# Perform online prediction using BOA There are 3 equivalent possibilities 1)
# start with an empty model and update the model sequentially
m1.BOA <- m0.BOA
for (i in 1:length(Y)) {
  m1.BOA <- predict(m1.BOA, newexperts = X[i, ], newY = Y[i], quiet = TRUE)
}

# 2) perform online prediction directly from the empty model
m2.BOA <- predict(m0.BOA, newexpert = X, newY = Y, online = TRUE, quiet = TRUE)

# 3) perform the online aggregation directly
m3.BOA <- mixture(Y = Y, experts = X, model = 'BOA', loss.type = 'square', quiet = TRUE)

# These predictions are equivalent:
identical(m1.BOA, m2.BOA) # TRUE
identical(m1.BOA, m3.BOA) # TRUE

# Display the results
summary(m3.BOA)
if(interactive()){
  plot(m1.BOA)
}

# Using d-dimensional time-series
#####

# Consider the above exemple of electricity consumption
# to be predicted every four weeks
YBlock <- seriesToBlock(X = Y, d = 4)
XBlock <- seriesToBlock(X = X, d = 4)

# The four-week-by-four-week predictions can then be obtained
# by directly using the `mixture` function as we did earlier.

```



```

MLpolBlock <- mixture(Y = YBlock, experts = XBlock, model = "MLpol", loss.type = "square",
                    quiet = TRUE)

# The predictions can finally be transformed back to a
# regular one dimensional time-series by using the function `blockToSeries`.

prediction <- blockToSeries(MLpolBlock$prediction)

#### Using the `online = FALSE` option

# Equivalent solution is to use the `online = FALSE` option in the predict function.
# The latter ensures that the model coefficients are not
# updated between the next four weeks to forecast.
MLpolBlock <- mixture(model = "BOA", loss.type = "square")
d = 4
n <- length(Y)/d
for (i in 0:(n-1)) {
  idx <- 4*i + 1:4 # next four weeks to be predicted
  MLpolBlock <- predict(MLpolBlock, newexperts = X[idx, ], newY = Y[idx], online = FALSE,
                      quiet = TRUE)
}

print(head(MLpolBlock$weights))

```

oracle

Compute oracle predictions

Description

The function `oracle` performs a strategie that cannot be defined online (in contrast to `mixture`). It requires in advance the knowledge of the whole data set Y and the expert advice to be well defined. Examples of oracles are the best fixed expert, the best fixed convex combination rule, the best linear combination rule, or the best expert that can shift a few times.

Usage

```

oracle(
  Y,
  experts,
  model = "convex",
  loss.type = "square",
  awake = NULL,
  lambda = NULL,
  niter = NULL,
  ...
)

```

Arguments

<code>Y</code>	A vector containing the observations to be predicted.
<code>experts</code>	A matrix containing the experts forecasts. Each column corresponds to the predictions proposed by an expert to predict Y . It has as many columns as there are experts.
<code>model</code>	A character string specifying the oracle to use or a list with a component name specifying the oracle and any additional parameter needed. Currently available oracles are: 'expert' The best fixed (constant over time) expert oracle. 'convex' The best fixed convex combination (vector of non-negative weights that sum to 1) 'linear' The best fixed linear combination of expert 'shifting' It computes for all number m of switches the sequence of experts with at most m shifts that would have performed the best to predict the sequence of observations in Y .
<code>loss.type</code>	character, list or function. character Name of the loss to be applied ('square', 'absolute', 'percentage', or 'pinball'); list When using pinball loss: list with field name equal to 'pinball' and field tau equal to the required quantile in $[0,1]$; function A custom loss as a function of two parameters.
<code>awake</code>	A matrix specifying the activation coefficients of the experts. Its entries lie in $[0,1]$. Possible if some experts are specialists and do not always form and suggest prediction. If the expert number k at instance t does not form any prediction of observation Y_t , we can put <code>awake[t,k]=0</code> so that the mixture does not consider expert k in the mixture to predict Y_t . Remark that to compute the best expert oracle, the performance of unactive (or partially active) experts is computed by using the prediction of the uniform average of active experts.
<code>lambda</code>	A positive number used by the 'linear' oracle only. A possible L_2 regularization parameter for computing the linear oracle (if the design matrix is not identifiable)
<code>niter</code>	A positive integer for 'convex' and 'linear' oracles if direct computation of the oracle is not implemented. It defines the number of optimization steps to perform in order to approximate the oracle (default value is 3).
<code>...</code>	Additional parameters that are passed to <code>optim</code> function in order to perform convex optimization (see parameter <code>niter</code>).

Value

An object of class 'oracle' that contains:

<code>loss</code>	The average loss suffered by the oracle. For the 'shifting' oracle, it is a vector of length T where T is the number of instance to be predicted (i.e., the length of the sequence Y). The value of <code>loss(m)</code> is the loss (determined by the parameter <code>loss.type</code>) suffered by the best sequence of expert with at most $m-1$ shifts.
-------------------	--

coefficients	Not for the 'shifting' oracle. A vector containing the best weight vector corresponding to the oracle.
prediction	Not for the 'shifting' oracle. A vector containing the predictions of the oracle.
rmse	If loss.type is the square loss (default) only. The root mean square error (i.e., it is the square root of loss).

Author(s)

Pierre Gaillard <pierre@gaillard.me>

plot.mixture	<i>Plot an object of class mixture</i>
--------------	--

Description

provides different diagnostic plots for an aggregation procedure.

Usage

```
## S3 method for class 'mixture'
plot(
  x,
  pause = FALSE,
  col = NULL,
  alpha = 0.01,
  dynamic = T,
  type = "all",
  max_experts = 50,
  col_by_weight = TRUE,
  ...
)
```

Arguments

x	an object of class mixture. If awake is provided (i.e., some experts are unactive), their residuals and cumulative losses are computed by using the predictions of the mixture.
pause	if set to TRUE (default) displays the plots separately, otherwise on a single page
col	the color to use to represent each experts, if set to NULL (default) use <code>RColorBrewer::brewer.pal(...)</code>
alpha	numeric. Smoothing parameter for contribution plot (parameter 'f' of function lowess).
dynamic	boolean. If TRUE, graphs are generated with <code>rAmCharts</code> , else with base R.
type	char. <ul style="list-style-type: none"> 'all' Display all the graphs ;

- 'plot_weight', 'boxplot_weight', 'dyn_avg_loss', 'cumul_res', 'avg_loss', 'contrib' Display the selected graph alone.

max_experts	integer. Maximum number of experts to be displayed (only the more influential).
col_by_weight	boolean. If TRUE (default), colors are ordered by weights of each expert, else by column
...	additional plotting parameters

Value

plots representing: plot of weights of each expert in function of time, boxplots of these weights, cumulative loss $L_T = \sum_{t=1}^T l_{i,t}$ of each expert in function of time, cumulative residuals $\sum_{t=1}^T (y_t - f_{i,t})$ of each expert's forecast in function of time, average loss suffered by the experts and the contribution of each expert to the aggregation $p_{i,t} f_{i,t}$ in function of time.

Author(s)

Pierre Gaillard <pierre@gaillard.me>

Yannig Goude <yannig.goude@edf.fr>

See Also

See [opera-package](#) and opera-vignette for a brief example about how to use the package.

plot.oracle	<i>Plot an aggregation procedure</i>
-------------	--------------------------------------

Description

oracle plot. It has one optional arguments.

Usage

```
## S3 method for class 'oracle'
plot(x, sort = TRUE, col = NULL, dynamic = TRUE, ...)
```

Arguments

x	An object of class oracle.
sort	if set to TRUE (default), it sorts the experts by performance before the plots.
col	colors
dynamic	If TRUE, graphs are generated with rAmCharts, else with base R.
...	additional arguments to function plot.

plot_ridge_weights *Functions to render dynamic mixture graphs using rAmCharts*

Description

Functions to render dynamic mixture graphs using rAmCharts

Usage

```
plot_ridge_weights(data, colors = NULL, max_experts = 50, round = 3)
plot_weights(data, colors = NULL, max_experts = 50, round = 3)
boxplot_weights(data, colors = NULL, max_experts = 50)
plot_dyn_avg_loss(data, colors = NULL, max_experts = 50, round = 3)
plot_cumul_res(data, colors = NULL, max_experts = 50, round = 3)
plot_avg_loss(data, colors = NULL, max_experts = 50, round = 3)
plot_contrib(data, colors = NULL, alpha = 0.1, max_experts = 50, round = 3)
```

Arguments

data	mixture object. Displays graphs.
colors	character. Colors of the lines and bullets.
max_experts	integer. Maximum number of experts to be displayed (only the more influential).
round	integer. Precision of the displayed values.
alpha	numeric. Smoothing parameter for contribution plot (parameter 'f' of function lowess).

Value

a rAmCharts plot

plt_oracle_convex *Functions to render dynamic oracle graphs using rAmCharts*

Description

Functions to render dynamic oracle graphs using rAmCharts

Usage

```
plt_oracle_convex(data, colors, round = 2)
```

Arguments

data	named vector. Vector of values to be displayed.
colors	character. Colors to be used.
round	integer (2). Precision of the values in the tooltips..

Value

a rAmCharts plot

predict.mixture *Predict method for Mixture models*

Description

Performs sequential predictions and updates of a mixture object based on new observations and expert advice.

Usage

```
## S3 method for class 'mixture'
predict(
  object,
  newexperts = NULL,
  newY = NULL,
  awake = NULL,
  online = TRUE,
  type = c("model", "response", "weights", "all"),
  use_cpp = getOption("opera_use_cpp", default = FALSE),
  quiet = TRUE,
  ...
)
```

Arguments

object	Object of class inheriting from 'mixture'
newexperts	An optional matrix in which to look for expert advice with which predict. If omitted, the past predictions of the object are returned and the object is not updated.
newY	An optional matrix with d columns (or vector if $d = 1$) of observations to be predicted. If provided, it should have the same number of rows as the number of rows of newexperts. If omitted, the object (i.e, the aggregation rule) is not updated.
awake	An optional array specifying the activation coefficients of the experts. It must have the same dimension as experts. Its entries lie in $[0, 1]$. Possible if some experts are specialists and do not always form and suggest prediction. If the expert number k at instance t does not form any prediction of observation Y_t , we can put <code>awake[t, k]=0</code> so that the mixture does not consider expert k in the mixture to predict Y_t .
online	A boolean determining if the observations in newY are predicted sequentially (by updating the object step by step) or not. If FALSE, the observations are predicting using the object (without using any past information in newY). If TRUE, newY and newexperts should not be null.
type	Type of prediction. It can be model return the updated version of object (using newY and newexperts). response return the forecasts. If type is 'model', forecasts can also be obtained from the last values of object\$prediction. weights return the weights assigned to the expert advice to produce the forecasts. If type is 'model', forecasts can also be obtained from the last rows of object\$weights. all return a list containing 'model', 'response', and 'weights'.
use_cpp	boolean. Whether or not to use cpp optimization to fasten the computations. This option is not yet compatible with the use of custom loss function.
quiet	boolean. Whether or not to display progress bars.
...	further arguments are ignored

Value

predict.mixture produces a matrix of predictions (type = 'response'), an updated object (type = 'model'), or a matrix of weights (type = 'weights').

Description

The functions `seriesToBlock` and `blockToSeries` convert 1-dimensional series into series of higher dimension. For instance, suppose you have a time-series that consists of $T = 100$ days of $d = 24$ hours. The function `seriesToBlock` converts the time-series X of $Td = 2400$ observations into a matrix of size $c(T=100, d=24)$, where each line corresponds to a specific day. This function is useful if you need to perform the prediction day by day, instead of hour by hour. The function can also be used to convert a matrix of expert prediction of dimension $c(dT, K)$ where K is the number of experts, into an array of dimension $c(T, d, K)$. The new arrays of observations and of expert predictions can be given to the aggregation rule procedure to perform d -dimensional predictions (i.e., day predictions).

Usage

```
seriesToBlock(X, d)
```

```
blockToSeries(X)
```

Arguments

X	An array or a vector to be converted.
d	A positive integer defining the block size.

Details

The function `blockToSeries` performs the inverse operation.

Index

- * **~models**
 - mixture, [10](#)
- * **~ts**
 - mixture, [10](#)
- * **datasets**
 - electric_load, [7](#)
- * **package**
 - opera-package, [2](#)

[auglag](#), [12](#)

[blockToSeries \(seriesToBlock\)](#), [23](#)
[boxplot_weights \(plot_ridge_weights\)](#), [21](#)

[check_loss](#), [5](#)
[check_matrix](#), [6](#)

[electric_load](#), [7](#)

[FTRL](#), [7](#), [12](#)

[loss](#), [9](#)
[lowess](#), [19](#), [21](#)

[mixture](#), [10](#), [17](#)

[opera \(opera-package\)](#), [2](#)
[opera-package](#), [2](#)
[optim](#), [18](#)
[oracle](#), [11](#), [17](#)

[plot.mixture](#), [19](#)
[plot.oracle](#), [20](#)
[plot_avg_loss \(plot_ridge_weights\)](#), [21](#)
[plot_contrib \(plot_ridge_weights\)](#), [21](#)
[plot_cumul_res \(plot_ridge_weights\)](#), [21](#)
[plot_dyn_avg_loss \(plot_ridge_weights\)](#),
[21](#)
[plot_ridge_weights](#), [21](#)
[plot_weights \(plot_ridge_weights\)](#), [21](#)
[plt_oracle_convex](#), [22](#)

[predict.mixture](#), [22](#)
[print.mixture \(mixture\)](#), [10](#)
[seriesToBlock](#), [23](#)
[summary.mixture \(mixture\)](#), [10](#)