

Package ‘BT’

January 20, 2025

Title (Adaptive) Boosting Trees Algorithm

Version 0.4

Date 2023-08-19

Author Gireg Willame [aut, cre, cph]

Maintainer Gireg Willame <gireg.willame@gmail.com>

Depends R (>= 4.0)

Imports rpart, stats, statmod, parallel

Suggests rmarkdown, knitr, testthat (>= 3.0.0)

Description Performs (Adaptive) Boosting Trees for Poisson distributed response variables, using log-link function.

The code approach is similar to the one used in 'gbm'/gbm3'. Moreover, each tree in the expansion is built thanks to the 'rpart' package.

This package is based on following books and articles

Denuit, M., Hainaut, D., Trufin, J. (2019) <doi:10.1007/978-3-030-25820-7>

Denuit, M., Hainaut, D., Trufin, J. (2019) <doi:10.1007/978-3-030-57556-4>

Denuit, M., Hainaut, D., Trufin, J. (2019) <doi:10.1007/978-3-030-25827-6>

Denuit, M., Hainaut, D., Trufin, J. (2022) <doi:10.1080/03461238.2022.2037016>

Denuit, M., Huyghe, J., Trufin, J. (2022) <https://dial.uclouvain.be/pr/boreal/fr/object/boreal%3A244325/datastream/PDF_01/view>

Denuit, M., Trufin, J., Verdebout, T. (2022) <<https://dial.uclouvain.be/pr/boreal/fr/object/boreal%3A268577>>.

URL <https://github.com/GiregWillame/BT/>

BugReports <https://github.com/GiregWillame/BT/issues/>

License GPL (>= 3)

Encoding UTF-8

RoxygenNote 7.2.3

Config/testthat/edition 3

VignetteBuilder knitr

LazyData true

NeedsCompilation no

Repository CRAN

Date/Publication 2023-08-19 09:22:33 UTC

Contents

BT	2
BTCVFit	6
BTFit	7
BT_call	9
BT_devTweedie	12
BT_more	13
BT_perf	14
BT_Simulated_Data	16
predict.BTFit	17
print.BTFit	18
summary.BTFit	19

Index **22**

BT *(Adaptive) Boosting Trees (ABT/BT) Algorithm.*

Description

Performs the (Adaptive) Boosting Trees algorithm. This code prepares the inputs and calls the function `BT_call`. Each tree in the process is built thanks to the `rpart` function. In case of cross-validation, this function prepares the folds and performs multiple calls to the fitting function `BT_call`.

Usage

```
BT(
  formula = formula(data),
  data = list(),
  tweedie.power = 1,
  ABT = TRUE,
  n.iter = 100,
  train.fraction = 1,
  interaction.depth = 4,
  shrinkage = 1,
  bag.fraction = 1,
  colsample.bytree = NULL,
  keep.data = TRUE,
  is.verbose = FALSE,
  cv.folds = 1,
  folds.id = NULL,
  n.cores = 1,
```

```

tree.control = rpart.control(xval = 0, maxdepth = (if (!is.null(interaction.depth)) {
  interaction.depth
} else {
  10
}), cp = -Inf, minsplit = 2),
weights = NULL,
seed = NULL,
...
)

```

Arguments

<code>formula</code>	a symbolic description of the model to be fit. Note that the offset isn't supported in this algorithm. Instead, everything is performed with a log-link function and a direct relationship exist between response, offset and weights.
<code>data</code>	an optional data frame containing the variables in the model. By default the variables are taken from <code>environment(formula)</code> , typically the environment from which BT is called. If <code>keep.data=TRUE</code> in the initial call to BT then BT stores a copy with the object (up to the variables used).
<code>tweedie.power</code>	Experimental parameter currently not used - Set to 1 referring to Poisson distribution.
<code>ABT</code>	a boolean parameter. If <code>ABT=TRUE</code> an adaptive boosting tree algorithm is built whereas if <code>ABT=FALSE</code> an usual boosting tree algorithm is run. By default, it is set to <code>TRUE</code> .
<code>n.iter</code>	the total number of iterations to fit. This is equivalent to the number of trees and the number of basis functions in the additive expansion. Please note that the initialization is not taken into account in the <code>n.iter</code> . More explicitly, a weighted average initializes the algorithm and then <code>n.iter</code> trees are built. Moreover, note that the <code>bag.fraction</code> , <code>colsample.bytree</code> , ... are not used for this initializing phase. By default, it is set to 100.
<code>train.fraction</code>	the first <code>train.fraction * nrow(data)</code> observations are used to fit the BT and the remainder are used for computing out-of-sample estimates (also known as validation error) of the loss function. By default, it is set to 1 meaning no out-of-sample estimates.
<code>interaction.depth</code>	the maximum depth of variable interactions: 1 builds an additive model, 2 builds a model with up to two-way interactions, etc. This parameter can also be interpreted as the maximum number of non-terminal nodes. By default, it is set to 4. Please note that if this parameter is <code>NULL</code> , all the trees in the expansion are built based on the <code>tree.control</code> parameter only, independently of the <code>ABT</code> value. This option is devoted to advanced users only and allows them to benefit from the full flexibility of the implemented algorithm.
<code>shrinkage</code>	a shrinkage parameter (in the interval (0,1]) applied to each tree in the expansion. Also known as the learning rate or step-size reduction. By default, it is set to 1.
<code>bag.fraction</code>	the fraction of independent training observations randomly selected to propose the next tree in the expansion. This introduces randomness into the model fit. If

`bag.fraction<1` then running the same model twice will result in similar but different fits. Please note that if this parameter is used the `BTErrors$training.error` corresponds to the normalized in-bag error and the out-of-bag improvements are computed and stored in `BTErrors$oob.improvement`. See [BTFit](#) for more details. By default, it is set to 1.

<code>colsample.bytree</code>	each tree will be trained on a random subset of <code>colsample.bytree</code> number of features. Each tree will consider a new random subset of features from the formula, adding variability to the algorithm and reducing computation time. <code>colsample.bytree</code> will be bounded between 1 and the number of features considered in the formula. By default, it is set to NULL meaning no effect.
<code>keep.data</code>	a boolean variable indicating whether to keep the data frames. This is particularly useful if one wants to keep track of the initial data frames and is further used for predicting in case any data frame is specified. Note that in case of cross-validation, if <code>keep.data=TRUE</code> the initial data frames are saved whereas the cross-validation samples are not. By default, it is set to FALSE.
<code>is.verbose</code>	if <code>is.verbose=TRUE</code> , the BT will print out the algorithm progress. By default, it is set to FALSE.
<code>cv.folds</code>	a positive integer representing the number of cross-validation folds to perform. If <code>cv.folds>1</code> then BT, in addition to the usual fit, will perform a cross-validation and calculate an estimate of generalization error returned in <code>BTErrors\$cv.error</code> . By default, it is set to 1 meaning no cross-validation.
<code>folds.id</code>	an optional vector of values identifying what fold each observation is in. If supplied, this parameter prevails over <code>cv.folds</code> . By default, <code>folds.id = NULL</code> meaning that no folds are defined.
<code>n.cores</code>	the number of cores to use for parallelization. This parameter is used during the cross-validation. This parameter is bounded between 1 and the maximum number of available cores. By default, it is set to 1 leading to a sequential approach.
<code>tree.control</code>	for advanced user only. It allows to define additional tree parameters that will be used at each iteration. See rpart.control for more information.
<code>weights</code>	optional vector of weights used in the fitting process. These weights must be positive but do not need to be normalized. By default, it is set to NULL which corresponds to an uniform weight of 1 for each observation.
<code>seed</code>	optional number used as seed. Please note that if <code>cv.folds>1</code> , the <code>parLapply</code> function is called. Therefore, the seed (if defined) used inside each fold will be a multiple of the seed parameter.
<code>...</code>	not currently used.

Details

The NA values are currently dropped using `na.omit`.

Value

a [BTFit](#) object.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the gbm3 package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.

M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BTFit](#), [BTCVFit](#), [BT_call](#), [BT_perf](#), [predict.BTFit](#), [summary.BTFit](#), [print.BTFit](#), [.BT_cv_errors](#).

Examples

```
## Load dataset.
dataset <- BT::BT_Simulated_Data

## Fit a Boosting Tree model.
BT_algo <- BT(formula = Y_normalized ~ Age + Sport + Split + Gender, # formula
              data = dataset, # data
              ABT = FALSE, # Classical Boosting Tree
              n.iter = 200,
              train.fraction = 0.8,
              interaction.depth = 3,
              shrinkage = 0.01,
              bag.fraction = 0.5,
              colsample.bytree = 2, # 2 explanatory variable used at each iteration.
              keep.data = FALSE, # Do not keep a data copy.
              is.verbose = FALSE, # Do not print progress.
              cv.folds = 3, # 3-cv will be performed.
              folds.id = NULL ,
              n.cores = 1,
              weights = ExpoR, # <=> Poisson model on response Y with ExpoR in offset.
              seed = NULL)
```

```

## Determine the model performance and plot results.
best_iter_val <- BT_perf(BT_algo, method='validation')
best_iter_oob <- BT_perf(BT_algo, method='OOB', oobag.curve = TRUE)
best_iter_cv <- BT_perf(BT_algo, method='cv', oobag.curve = TRUE)

best_iter <- best_iter_val

## Variable influence and plot results.
# Based on the first iteration.
variable_influence1 <- summary(BT_algo, n.iter = 1)
# Using all iterations up to best_iter.
variable_influence_best_iter <- summary(BT_algo, n.iter = best_iter)

## Print results : call, best_iters and summarized relative influence.
print(BT_algo)

## Model predictions.
# Predict on the link scale, using only the best_iter tree.
pred_single_iter <- predict(BT_algo, newdata = dataset,
                           n.iter = best_iter, type = 'link', single.iter = TRUE)
# Predict on the response scale, using the first best_iter.
pred_best_iter <- predict(BT_algo, newdata = dataset,
                          n.iter = best_iter, type = 'response')

```

BTCVFit

BTCVFit

Description

These are objects representing CV fitted boosting trees.

Details

CV (Adaptive) Boosting Tree Model Object.

Value

a list of [BTFit](#) objects with each element corresponding to a specific BT fit on a particular fold

Structure

The following components must be included in a legitimate BTCVFit object.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the gbm3 package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.
- M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.
- M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BT](#).

BTFit

BTFit

Description

These are objects representing fitted boosting trees.

Details

Boosting Tree Model Object.

Value

BTInit	an object of class BTInit containing the initial fitted value <code>initFit</code> , the initial <code>training.error</code> and the initial <code>validation.error</code> if any.
BTErrors	an object of class BTErrors containing the vectors of errors for each iteration performed (excl. the initialization). More precisely, it contains the <code>training.error</code> , <code>validation.error</code> if <code>train.fraction</code> <1 and the <code>oob.improvement</code> if <code>bag.fraction</code> < 1. Moreover, if a cross-validation approach was performed, a vector of cross-validation errors <code>cv.error</code> as a function of boosting iteration is also stored in this object.
BTIndivFits	an object of class BTIndivFits containing the list of each individual tree fitted at each boosting iteration.
distribution	the Tweedie power (and so the distribution) that has been used to perform the algorithm. It will currently always output 1.

<code>var.names</code>	a vector containing the names of the explanatory variables.
<code>response</code>	the name of the target/response variable.
<code>w</code>	a vector containing the weights used.
<code>seed</code>	the used seed, if any.
<code>BTData</code>	if <code>keep.data=TRUE</code> , an object of class <code>BTData</code> containing the <code>training.set</code> and <code>validation.set</code> (can be <code>NULL</code> if not used). These data frames are reduced to the used variables, that are the response and explanatory variables. Note that in case of cross-validation, even if <code>keep.data=TRUE</code> the folds will not be kept. In fact, only the data frames related to the original fit (i.e. on the whole training set) will be saved.
<code>BTParams</code>	an object of class <code>BTParams</code> containing all the (Adaptive) boosting tree parameters. More precisely, it contains the <code>ABT</code> , <code>train.fraction</code> , <code>shrinkage</code> , <code>interaction.depth</code> , <code>bag.fraction</code> , <code>n.iter</code> , <code>colsample.bytree</code> and <code>tree.control</code> parameter values.
<code>keep.data</code>	the <code>keep.data</code> parameter value.
<code>is.verbose</code>	the <code>is.verbose</code> parameter value.
<code>fitted.values</code>	the training set fitted values on the score scale using all the <code>n.iter</code> (and initialization) iterations.
<code>cv.folds</code>	the number of cross-validation folds. Set to 1 if no cross-validation performed.
<code>call</code>	the original call to the BT algorithm.
<code>Terms</code>	the <code>model.frame</code> terms argument.
<code>folds</code>	a vector of values identifying to which fold each observation is in. This argument is not present if there is no cross-validation. On the other hand, it corresponds to <code>folds.id</code> if it was initially defined by the user.
<code>cv.fitted</code>	a vector containing the cross-validation fitted values, if a cross-validation was performed. More precisely, for a given observation, the prediction will be furnished by the <code>cv-model</code> for which this specific observation was out-of-fold. See predict.BTCVFit for more details.

Structure

The following components must be included in a legitimate `BTFit` object.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the `gbm3` package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.

M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BT](#).

BT_call

(Adaptive) Boosting Trees (ABT/BT) fit.

Description

Fit a (Adaptive) Boosting Trees algorithm. This is for "power" users who have a large number of variables and wish to avoid calling `model.frame` which can be slow in this instance. This function is in particular called by [BT](#). It is mainly split in two parts, the first one considers the initialization (see `BT_callInit`) whereas the second performs all the boosting iterations (see `BT_callBoosting`). By default, this function does not perform input checks (those are all done in [BT](#)) and all the parameters should be given in the right format. We therefore suppose that the user is aware of all the choices made.

Usage

```
BT_call(  
  training.set,  
  validation.set,  
  tweedie.power,  
  respVar,  
  w,  
  explVar,  
  ABT,  
  tree.control,  
  train.fraction,  
  interaction.depth,  
  bag.fraction,  
  shrinkage,  
  n.iter,  
  colsample.bytree,  
  keep.data,  
  is.verbose
```

```

)

BT_callInit(training.set, validation.set, tweedie.power, respVar, w)

BT_callBoosting(
  training.set,
  validation.set,
  tweedie.power,
  ABT,
  tree.control,
  interaction.depth,
  bag.fraction,
  shrinkage,
  n.iter,
  colsample.bytree,
  train.fraction,
  keep.data,
  is.verbose,
  respVar,
  w,
  explVar
)

```

Arguments

training.set	a data frame containing all the related variables on which one wants to fit the algorithm.
validation.set	a held-out data frame containing all the related variables on which one wants to assess the algorithm performance. This can be NULL.
tweedie.power	Experimental parameter currently not used - Set to 1 referring to Poisson distribution.
respVar	the name of the target/response variable.
w	a vector of weights.
explVar	a vector containing the name of explanatory variables.
ABT	a boolean parameter. If ABT=TRUE an adaptive boosting tree algorithm is built whereas if ABT=FALSE an usual boosting tree algorithm is run.
tree.control	allows to define additional tree parameters that will be used at each iteration. See rpart.control for more information.
train.fraction	the first <code>train.fraction * nrows(data)</code> observations are used to fit the BT and the remainder are used for computing out-of-sample estimates (also known as validation error) of the loss function. It is mainly used to report the value in the BTfit object.
interaction.depth	the maximum depth of variable interactions: 1 builds an additive model, 2 builds a model with up to two-way interactions, etc. This parameter can also be interpreted as the maximum number of non-terminal nodes. By default, it is set to 4.

Please note that if this parameter is NULL, all the trees in the expansion are built based on the `tree.control` parameter only. This option is devoted to advanced users only and allows them to benefit from the full flexibility of the implemented algorithm.

<code>bag.fraction</code>	the fraction of independent training observations randomly selected to propose the next tree in the expansion. This introduces randomness into the model fit. If <code>bag.fraction < 1</code> then running the same model twice will result in similar but different fits. BT uses the R random number generator, so <code>set.seed</code> ensures the same model can be reconstructed. Please note that if this parameter is used the <code>BTErrors\$training.error</code> corresponds to the normalized in-bag error.
<code>shrinkage</code>	a shrinkage parameter applied to each tree in the expansion. Also known as the learning rate or step-size reduction.
<code>n.iter</code>	the total number of iterations to fit. This is equivalent to the number of trees and the number of basis functions in the additive expansion. Please note that the initialization is not taken into account in the <code>n.iter</code> . More explicitly, a weighted average initializes the algorithm and then <code>n.iter</code> trees are built. Moreover, note that the <code>bag.fraction</code> , <code>colsample.bytree</code> , ... are not used for this initializing phase.
<code>colsample.bytree</code>	each tree will be trained on a random subset of <code>colsample.bytree</code> number of features. Each tree will consider a new random subset of features from the formula, adding variability to the algorithm and reducing computation time. <code>colsample.bytree</code> will be bounded between 1 and the number of features considered.
<code>keep.data</code>	a boolean variable indicating whether to keep the data frames. This is particularly useful if one wants to keep track of the initial data frames and is further used for predicting in case any data frame is specified. Note that in case of cross-validation, if <code>keep.data=TRUE</code> the initial data frames are saved whereas the cross-validation samples are not.
<code>is.verbose</code>	if <code>is.verbose=TRUE</code> , the BT will print out the algorithm progress.

Value

a `BTFit` object.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the `gbm3` package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.

M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BTFit](#), [BTCVFit](#), [BT_perf](#), [predict.BTFit](#), [summary.BTFit](#), [print.BTFit](#), [.BT_cv_errors](#).

BT_devTweedie

Deviance function for the Tweedie family.

Description

Compute the deviance for the Tweedie family case.

Usage

```
BT_devTweedie(y, mu, tweedieVal, w = NULL)
```

Arguments

y	a vector containing the observed values.
mu	a vector containing the fitted values.
tweedieVal	a numeric representing the Tweedie Power. It has to be a positive number outside of the interval]0,1[.
w	an optional vector of weights.

Details

This function computes the Tweedie related deviance. The latter is defined as:

$$\begin{aligned}
 d(y, mu, w) &= w(y - mu)^2, \text{ if } tweedieVal = 0; \\
 d(y, mu, w) &= 2w(y \log(y/mu) + mu - y), \text{ if } tweedieVal = 1; \\
 d(y, mu, w) &= 2w(\log(mu/y) + y/mu - 1), \text{ if } tweedieVal = 2; \\
 d(y, mu, w) &= 2w(\max(y, 0)^{(2-p)/((1-p)(2-p))} - ymu^{(1-p)/(1-p)} + mu^{(2-p)/(2-p)}), \text{ else.}
 \end{aligned}$$

Value

A vector of individual deviance contribution.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the gbm3 package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.

M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BT](#), [BT_call](#).

BT_more

Perform additional boosting iterations.

Description

Method to perform additional iterations of the Boosting Tree algorithm, starting from an initial [BTFit](#) object. This does not support further cross-validation. Moreover, this approach is only allowed if `keep.data=TRUE` in the original call.

Usage

```
BT_more(BTFit_object, new.n.iter = 100, is.verbose = FALSE, seed = NULL)
```

Arguments

BTFit_object	a BTFit object.
new.n.iter	number of new boosting iterations to perform.
is.verbose	a logical specifying whether or not the additional fitting should run "noisely" with feedback on progress provided to the user.
seed	optional seed used to perform the new iterations. By default, no seed is set.

Value

Returns a new `BTFit` object containing the initial call as well as the new iterations performed.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the `gbm3` package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.

M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

`BT`, `BTFit`.

BT_perf

Performance assessment.

Description

Function to compute the performances of a fitted boosting tree.

Usage

```
BT_perf(  
  BTFit_object,  
  plot.it = TRUE,  
  oobag.curve = FALSE,  
  overlay = TRUE,  
  method,  
  main = ""  
)
```

Arguments

BTFit_object	a BTFit object resulting from an initial call to BT
plot.it	a boolean indicating whether to plot the performance measure. Setting <code>plot.it = TRUE</code> creates two plots. The first one plots the <code>object\$BTErrors\$training.error</code> (in black) as well as the <code>object\$BTErrors\$validation.error</code> (in red) and/or the <code>object\$BTErrors\$cv.error</code> (in green) depending on the method and parametrization. These values are plotted as a function of the iteration number. The scale of the error measurement, shown on the left vertical axis, depends on the arguments used in the initial call to BT and the chosen method.
oobag.curve	indicates whether to plot the out-of-bag performance measures in a second plot. Note that this option makes sense if the <code>bag.fraction</code> was properly defined in the initial call to BT .
overlay	if set to <code>TRUE</code> and <code>oobag.curve=TRUE</code> then a right y-axis is added and the estimated cumulative improvement in the loss function is plotted versus the iteration number.
method	indicates the method used to estimate the optimal number of boosting iterations. Setting <code>method = "OOB"</code> computes the out-of-bag estimate and <code>method = "validation"</code> uses the validation dataset to compute an out-of-sample estimate. Finally, setting <code>method = "cv"</code> extracts the optimal number of iterations using cross-validation, if BT was called with <code>cv.folds > 1</code> . If missing, a guessing method is applied.
main	optional parameter that allows the user to define specific plot title.

Value

Returns the estimated optimal number of iterations. The method of computation depends on the `method` argument.

Author(s)

Gireg Willame <g.willame@detralytics.eu>

This package is inspired by the `gbm3` package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.

M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BT](#), [BT_call](#).

BT_Simulated_Data *Simulated Database*.

Description

A simulated database used for examples and vignettes. The variables are related to a motor insurance pricing context.

Usage

```
BT_Simulated_Data
```

Format

A simulated data frame with 50,000 rows and 7 columns, containing simulation of different policyholders:

Gender Gender, varying between male and female.

Age Age, varying from 18 to 65years old.

Split Noisy variable, not used to simulate the response variable. It allows to assess how the algorithm handle these features.

Sport Car type, varying between yes (sport car) or no.

ExpoR Yearly exposure-to-risk, varying between 0 and 1.

Y Yearly claim number, simulated thanks to Poisson distribution.

Y_normalized Yearly claim frequency, corresponding to the ratio between Y and ExpoR.

predict.BTFit	<i>Predict method for BT Model fits.</i>
---------------	--

Description

Predicted values based on a boosting tree model object.

Usage

```
## S3 method for class 'BTFit'  
predict(object, newdata, n.iter, type = "link", single.iter = FALSE, ...)
```

Arguments

object	a <code>BTFit</code> object.
newdata	data frame of observations for which to make predictions. If missing or not a data frame, if <code>keep.data=TRUE</code> in the initial fit then the original training set will be used.
n.iter	number of boosting iterations used for the prediction. This parameter can be a vector in which case predictions are returned for each iteration specified.
type	the scale on which the BT makes the predictions. Can either be "link" or "response". Note that, by construction, a log-link function is used during the fit.
single.iter	if <code>single.iter=TRUE</code> then <code>predict.BTFit</code> returns the predictions from the single tree <code>n.iter</code> .
...	not currently used.

Details

`predict.BTFit` produces a predicted values for each observation in `newdata` using the first `n.iter` boosting iterations. If `n.iter` is a vector then the result is a matrix with each column corresponding to the BT predictions with `n.iter[1]` boosting iterations, `n.iter[2]` boosting iterations, and so on.

As for the fit, the predictions do not include any offset term. In the Poisson case, please remind that a weighted approach is initially favored.

Value

Returns a vector of predictions. By default, the predictions are on the score scale. If `type = "response"`, then BT converts back to the same scale as the outcome. Note that, a log-link is supposed by construction.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the `gbm3` package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.
- M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.
- M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BT](#), [BTFit](#).

print.BTFit

Printing function.

Description

Function to print the BT results.

Usage

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

Arguments

x a [BTFit](#) object.
... arguments passed to `print.default`.

Details

Print the different input parameters as well as obtained results (best iteration/performance & relative influence) given the chosen approach.

Value

No value returned.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the gbm3 package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.

M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.

M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BT](#), [.BT_relative_influence](#), [BT_perf](#).

summary.BTFit

Summary of a BTFit object.

Description

Computes the relative influence of each variable in the BTFit object.

Usage

```
## S3 method for class 'BTFit'
summary(
  object,
  cBars = length(object$var.names),
  n.iter = object$BTParams$n.iter,
  plot_it = TRUE,
  order_it = TRUE,
  method = .BT_relative_influence,
  normalize = TRUE,
  ...
)
```

Arguments

object	a <code>BTFit</code> object.
cBars	the number of bars to plot. If <code>order=TRUE</code> only the variables with the cBars largest relative influence will appear in the barplot. If <code>order=FALSE</code> then the first cBars variables will appear in the barplot.
n.iter	the number of trees used to compute the relative influence. Only the first n.iter trees will be used.
plot_it	an indicator as to whether the plot is generated.
order_it	an indicator as to whether the plotted and/or returned relative influences are sorted.
method	the function used to compute the relative influence. Currently, only <code>.BT_relative_influence</code> is available (default value as well).
normalize	if TRUE returns the normalized relative influence.
...	additional argument passed to the plot function.

Details

Please note that the relative influence for variables having an original **negative** relative influence is forced to 0.

Value

Returns a data frame where the first component is the variable name and the second one is the computed relative influence, normalized to sum up to 100. Depending on the `plot_it` value, the relative influence plot will be performed.

Author(s)

Gireg Willame <gireg.willame@gmail.com>

This package is inspired by the `gbm3` package. For more details, see <https://github.com/gbm-developers/gbm3/>.

References

- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries II: Tree-Based Methods and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2019). **Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions**, *Springer Actuarial*.
- M. Denuit, D. Hainaut and J. Trufin (2022). **Response versus gradient boosting trees, GLMs and neural networks under Tweedie loss and log-link**. Accepted for publication in *Scandinavian Actuarial Journal*.
- M. Denuit, J. Huyghe and J. Trufin (2022). **Boosting cost-complexity pruned trees on Tweedie responses: The ABT machine for insurance ratemaking**. Paper submitted for publication.

M. Denuit, J. Trufin and T. Verdebout (2022). **Boosting on the responses with Tweedie loss functions**. Paper submitted for publication.

See Also

[BT](#), [.BT_relative_influence](#).

Index

* datasets

BT_Simulated_Data, 16

* methods

BTCVFit, 6

BTFit, 7

.BT_cv_errors, 5, 12

.BT_relative_influence, 19–21

BT, 2, 7, 9, 13–16, 18, 19, 21

BT_call, 2, 5, 9, 13, 16

BT_callBoosting (BT_call), 9

BT_callInit (BT_call), 9

BT_devTweedie, 12

BT_more, 13

BT_perf, 5, 12, 14, 19

BT_Simulated_Data, 16

BTCVFit, 5, 6, 12

BTFit, 4–6, 7, 11–15, 17, 18, 20

predict.BTCVFit, 8

predict.BTFit, 5, 12, 17

print.BTFit, 5, 12, 18

rpart, 2

rpart.control, 4, 10

summary.BTFit, 5, 12, 19