

Package ‘BartMixVs’

October 12, 2022

Type Package

Title Variable Selection Using Bayesian Additive Regression Trees

Version 1.0.0

Date 2022-05-02

Author Chuji Luo [aut, cre],
Michael J. Daniels [aut],
Robert McCulloch [ctb],
Rodney Sparapani [ctb]

Maintainer Chuji Luo <cjluo@uf1.edu>

Description

Bayesian additive regression trees (BART) provides flexible non-parametric modeling of mixed-type predictors for continuous and binary responses. This package is built upon CRAN R package 'BART', version 2.7 (<<https://github.com/cran/BART>>). It implements the three proposed variable selection approaches in the paper: Luo, C and Daniels, M. J. (2021), "Variable Selection Using Bayesian Additive Regression Trees." <[arXiv:2112.13998](https://arxiv.org/abs/2112.13998)>, and other three existing BART-based variable selection approaches.

License GPL (>= 2)

Depends R (>= 2.10), nlme, nnet

Imports Rcpp (>= 1.0.6), loo, mvtnorm

LinkingTo Rcpp

RoxygenNote 7.1.1

Encoding UTF-8

NeedsCompilation yes

Repository CRAN

Date/Publication 2022-05-05 05:50:02 UTC

R topics documented:

BartMixVs-package	2
abc.vs	3
bartModelMatrix	7

checkerboard	9
friedman	10
mc.abc.vs	11
mc.backward.vs	16
mc.cores.openmp	19
mc.pbart	20
mc.permute.vs	25
mc.pwbart	30
mc.wbart	32
medianInclusion.vs	37
mixone	41
mixtwo	42
pbart	43
permute.vs	48
predict.pbart	53
predict.wbart	54
pwbart	56
wbart	57

Index	64
--------------	-----------

BartMixVs-package	<i>Variable Selection Using Bayesian Additive Regression Trees</i>
-------------------	--

Description

This package provides implementations of existing BART-based variable selection approaches.

Details

Bayesian additive regression trees (BART) provides flexible nonparametric modeling of mixed-type predictors for continuous and binary responses. This package is built upon CRAN R package BART, version 2.7 (<https://github.com/cran/BART>). It implements the three proposed variable selection approaches of the paper, Luo, C and Daniels, MJ (2022), "Variable Selection Using Bayesian Additive Regression Trees", and other three existing BART-based variable selection approaches.

Author(s)

Chuji Luo <cjluo@ufl.edu>, Michael J. Daniels <daniels@ufl.edu>

Maintainer: Chuji Luo <cjluo@ufl.edu>

References

- LUO, C and DANIELS, MJ (2022). Variable Selection Using Bayesian Additive Regression Trees.
- BLEICH, J., KAPELNER, A., GEORGE, E. I. and JENSEN, S. T. (2014). Variable selection for BART: an application to gene regulation. *Ann. Appl. Stat.* 8 1750–1781.
- LINERO, A. R. (2018). Bayesian regression trees for high-dimensional prediction and variable selection. *J. Amer. Statist. Assoc.* 113 626– 636.

LIU, Y., ROCKOVÁ, V. and WANG, Y. (2021). Variable selection with ABC Bayesian forests. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* 83 453–481.

SPARAPANI, R., SPANBAUER, C. and MCCULLOCH, R. (2021). Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package. *J. Stat. Softw.* 97 1–66.

See Also

Optional links to other man pages

Examples

```
## Not run:
## Optional simple examples of the most important functions
## These can be in \dontrun{} and \donttest{} blocks.

## End(Not run)
```

abc.vs

Variable selection with ABC Bayesian forest

Description

This function implements the variable selection approach proposed in Liu, Rockova and Wang (2021). Rockova and Pas (2020) introduce a spike-and-forest prior which wraps the BART prior with a spike-and-slab prior on the model space. Due to intractable marginal likelihood, Liu, Rockova and Wang (2021) propose an approximate Bayesian computation (ABC) sampling method based on data-splitting to help sample from the model space with higher ABC acceptance rate.

Usage

```
abc.vs(
  x,
  y,
  nabc = 1000,
  tolerance = 0.1,
  threshold = 0.25,
  beta.params = c(1, 1),
  beta.theta = NA,
  split.ratio = 0.5,
  probit = FALSE,
  true.idx = NULL,
  analysis = TRUE,
  sparse = FALSE,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
```

```

cont = FALSE,
rm.const = TRUE,
k = 2,
power = 2,
base = 0.95,
split.prob = "polynomial",
ntree = 10L,
ndpost = 1,
nskip = 200,
keepevery = 1L,
printevery = 100L,
verbose = FALSE
)

```

Arguments

x	A matrix or a data frame of predictors values with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
y	A vector of response (continuous or binary) values.
nabc	The number of ABC samples, i.e., the number of subsets sampled from the model space.
tolerance	A number between 0 and 1; the nabc subsets are ranked by MSE in ascending order if the response variable is continuous (or by mean log loss (MLL) if the response variable is binary), and the top $\text{tolerance} \times 100\%$ of the subsets are accepted by ABC for selection.
threshold	A number between 0 and 1; within the ABC accepted subsets, predictors with MPVIP exceeding threshold are selected.
beta.params	A vector with two positive numbers; the spike-and-slab prior on the model space is assumed to be a beta-binomial prior, i.e., $\theta \text{Beta}(\text{beta.params}[1], \text{beta.params}[2])$ and each predictor is included into a model by $\text{Bernoulli}(\theta)$; only used when $\text{beta.theta}=\text{NA}$.
beta.theta	A number between 0 and 1; the probability that a predictor is included into a model; if $\text{beta.theta}=\text{NA}$, it is sampled from $\text{Beta}(\text{beta.params}[1], \text{beta.params}[2])$.
split.ratio	A number between 0 and 1; the data set (x, y) is split into a training set and a testing set according to the <code>split.ratio</code> .
probit	A Boolean argument indicating whether the response variable is binary or continuous; <code>probit=FALSE</code> (by default) means that the response variable is continuous.
true.idx	(Optional) A vector of indices of the true relevant predictors; if <code>true.idx</code> is provided and <code>analysis=TRUE</code> , metrics including precision, recall and F1 score are returned.
analysis	A Boolean argument indicating whether to perform variable selection; if <code>analysis=TRUE</code> , the best model selected by ABC Bayesian forest is returned; if <code>analysis=FALSE</code> , only return the visited models and their corresponding model errors.
sparse	A Boolean argument indicating whether to perform DART or BART.

xinfo	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
numcut	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in x . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of x . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of $x - 1$)</code> cut-point values are used.
usequants	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
cont	A Boolean argument indicating whether to assume all predictors are continuous.
rm.const	A Boolean argument indicating whether to remove constant predictors.
k	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (y) is internally scaled to the range from $-.5$ to $.5$. The bigger k is, the more conservative the fitting will be.
power	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
base	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
split.prob	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
ntree	The number of trees in the ensemble.
ndpost	The number of posterior samples returned.
nskip	The number of posterior samples burned in.
keepevery	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
printevery	As the MCMC runs, a message is printed every <code>printevery</code> iterations.
verbose	A Boolean argument indicating whether any messages are printed out.

Details

At each iteration of the algorithm, the data set is randomly split into a training set and a testing set according to a certain split ratio. The algorithm proceeds by sampling a subset from the spike-and-slab prior on the model space, fitting a BART model on the training set only with the predictors in the subset, and computing the root mean squared errors (RMSE) for the test set based on a posterior sample from the fitted BART model. Only those subsets that result in a low RMSE on the test set are kept for selection. ABC Bayesian forest selects predictors based on their marginal posterior variable inclusion probabilities (MPVIPs) which are estimated by computing the proportion of ABC accepted BART posterior samples that use the predictor at least one time. Given the MPVIPs, predictors with MPVIP exceeding a pre-specified threshold are selected.

See Liu, Rockova and Wang (2021) or Section 2.2.4 in Luo and Daniels (2021) for details.

Value

The function `abc.vs()` returns a list with the following components.

<code>theta</code>	The probability that a predictor is included into a model.
<code>models</code>	A matrix with <code>nabc</code> rows and <code>ncol(x)</code> columns; each row corresponds to a ABC model (or subset); if the (i, j) -th element is 1, it means that the j -th predictor is included in the i -th ABC model; if the (i, j) -th element is 0, it means that the j -th predictor is not included in the i -th ABC model.
<code>actual.models</code>	A matrix with <code>nabc</code> rows and <code>ncol(x)</code> columns; each row corresponds to a ABC BART posterior sample; if the (i, j) -th element is 1, it means that the j -th predictor is used as a split variable at least one time in the BART posterior sample of the i -th ABC model; if the (i, j) -th element is 0, it means that the j -th predictor is not used as a split variable in the BART posterior sample of the i -th ABC model.
<code>model.errors</code>	The vector of MSEs (or MLLs if the response variable is binary) for the <code>nabc</code> ABC models.
<code>idx</code>	The vector of indices (in terms of the row numbers of <code>models</code>) of the ABC accepted models which are the top <code>tolerance*100%</code> of the <code>nabc</code> ABC models when ranked by MSE or MLL in ascending order; only returned when <code>analysis=TRUE</code> .
<code>top.models</code>	A matrix with <code>length(idx)</code> rows and <code>ncol(x)</code> columns, representing the ABC accepted models; <code>top.models=models[idx,]</code> ; only returned when <code>analysis=TRUE</code> .
<code>top.actual.models</code>	A matrix with <code>length(idx)</code> rows and <code>ncol(x)</code> columns, representing the ABC accepted BART posterior samples; <code>top.models=actual.models[idx,]</code> ; only returned when <code>analysis=TRUE</code> .
<code>mip</code>	The vector of marginal posterior variable inclusion probabilities; only returned when <code>analysis=TRUE</code> .
<code>best.model</code>	The vector of predictors selected by ABC Bayesian forest; only returned when <code>analysis=TRUE</code> .
<code>precision</code>	The precision score for the ABC Bayesian forest; only returned when <code>analysis=TRUE</code> and <code>true.idx</code> is provided.
<code>recall</code>	The recall score for the ABC Bayesian forest; only returned when <code>analysis=TRUE</code> and <code>true.idx</code> is provided.
<code>f1</code>	The F1 score for the ABC Bayesian forest; only returned when <code>analysis=TRUE</code> and <code>true.idx</code> is provided.

Author(s)

Chuji Luo: <cjluo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.

Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.

Liu, Yi, Veronika Rockova, and Yuexi Wang (2021). "Variable selection with ABC Bayesian forests." *J. R. Stat. Soc. Ser. B. Stat. Methodol.* 83.3, pp. 453–481.

Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.

Rockova Veronika and Stephanie van der Pas (2020). "Posterior concentration for Bayesian regression trees and forests." *Ann. Statist.* 48.4, pp. 2108–2131.

See Also

[permute.vs](#), [medianInclusion.vs](#) and [mc.backward.vs](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## test abc.vs() function
res = abc.vs(data$X, data$Y, nabc=100, tolerance=0.1, threshold=0.25, beta.params=c(1.0, 1.0),
split.ratio=0.5, probit=FALSE, true.idx=c(1,2,6:8), ntree=10, ndpost=1, nskip=200, analysis=TRUE)
```

bartModelMatrix

Create a matrix out of a vector or data frame

Description

The external BART functions (e.g. `wbart()`) operate on matrices in memory. Therefore, if the user submits a vector or data frame, then this function converts it to a matrix. Also, it determines the number of cut points necessary for each column when asked to do so. This function is inherited from the CRAN package 'BART'.

Usage

```
bartModelMatrix(
  X,
  numcut = 0L,
  usequants = FALSE,
  type = 7,
  rm.const = FALSE,
  cont = FALSE,
  xinfo = NULL
)
```

Arguments

X	A vector or data frame where the matrix is created.
numcut	The maximum number of cut points to consider. If numcut=0, then return a matrix; otherwise, return a list containing a matrix X, a vector numcut and a list xinfo.
usequants	A Boolean argument indicating the way to generate cut points. If usequants=FALSE, then the cut points in xinfo are generated uniformly; otherwise, the quantiles are used for the cut points.
type	An integer between 1 and 9 determining which algorithm is employed in the function quantile().
rm.const	A Boolean argument indicating whether to remove constant variables.
cont	A Boolean argument indicating whether to assume all variables are continuous.
xinfo	A list (matrix) where the items (rows) are the predictors and the contents (columns) of the items are the cut points. If xinfo=NULL, BART will choose xinfo for the user.

Value

The function `bartModelMatrix()` returns a list with the following components.

X	A matrix with rows corresponding to observations and columns corresponding to predictors (after dummification).
numcut	A vector of <code>ncol(X)</code> integers with each indicating the number of cut points for the corresponding predictor.
rm.const	A vector of indicators for the predictors (after dummification) used in BART; when the indicator is negative, it refers to remove that predictor.
xinfo	A list (matrix) where the items (rows) are the predictors and the contents (columns) of the items are the cut points.
grp	A vector of group indices for predictors. For example, if 2 appears 3 times in <code>grp</code> , the second predictor of X is a categorical predictor with 3 levels.

Author(s)

Chuji Luo: <cjluo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.

Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[wbart](#) and [pbart](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## test bartModelMatrix() function
res = bartModelMatrix(data$X, numcut=100, usequants=FALSE, cont=FALSE, rm.const=TRUE)
```

checkerboard

Generate data for an example of Zhu, Zeng and Kosorok (2015)

Description

Generate data including responses and predictors values according to an example of Zhu, R., Zeng, D. and Kosorok, M. R. (2015). "Reinforcement learning trees." *J. Amer. Statist. Assoc.* **110** 1770–1784.

Usage

```
checkerboard(n, p, sigma, binary)
```

Arguments

n	The number of observations.
p	The number of predictors.
sigma	The error variance.
binary	A boolean argument: binary = TRUE indicates that binary responses are generated and binary = FALSE indicates that continuous responses are generated.

Details

Sample the predictors x_1, \dots, x_p from $\text{Normal}(0, \Sigma)$ with $\Sigma_{jk} = 0.3^{|j-k|}$, $j, k = 1, \dots, p$. If binary = FALSE, sample the continuous response y from $\text{Normal}(f_0(x), \sigma^2)$, where

$$f_0(x) = 2x_1 * x_4 + 2x_7 * x_{10}.$$

If binary = TRUE, sample the binary response y from $\text{Bernoulli}(\Phi(f_0(x)))$ where f_0 is defined above and Φ is the cumulative density function of the standard normal distribution.

Value

Return a list with the following components.

X	An n by p data frame representing predictors values, with each row corresponding an observation.
Y	A vector of length n representing response values.
f0	A vector of length n representing the values of $f_0(x)$.
sigma	The error variance which is only returned when binary = FALSE.
prob	A vector of length n representing the values of $\Phi(f_0(x))$, which is only returned when binary = TRUE.

Author(s)

Chuji Luo: <cjluo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Zhu, R., Zeng, D. and Kosorok, M. R. (2015). "Reinforcement learning trees." *J. Amer. Statist. Assoc.* **110** 1770–1784.

Examples

```
data = checkerboard(100, 10, 1, FALSE)
```

friedman

Generate data for an example of Friedman (1991)

Description

Generate data including responses and predictors values according to an example of Friedman, J. H. (1991). "Multivariate adaptive regression splines." *Ann. Statist.* **19** 1–141.

Usage

```
friedman(n, p, sigma, binary)
```

Arguments

n	The number of observations.
p	The number of predictors.
sigma	The error variance.
binary	A boolean argument: binary = TRUE indicates that binary responses are generated and binary = FALSE indicates that continuous responses are generated.

Details

Sample the predictors x_1, \dots, x_p from $\text{Uniform}(0, 1)$ independently. If `binary = FALSE`, sample the continuous response y from $\text{Normal}(f_0(x), \sigma^2)$, where

$$f_0(x) = 10\sin(\pi x_1 * x_2) + 20(x_3 - 0.5)^2 + 10x_4 + 5x_5.$$

If `binary = TRUE`, sample the binary response y from $\text{Bernoulli}(\Phi(f_0(x)))$ where f_0 is defined above and Φ is the cumulative density function of the standard normal distribution.

Value

Return a list with the following components.

X	An n by p data frame representing predictors values, with each row corresponding an observation.
Y	A vector of length n representing response values.
f0	A vector of length n representing the values of $f_0(x)$.
sigma	The error variance which is only returned when <code>binary = FALSE</code> .
prob	A vector of length n representing the values of $\Phi(f_0(x))$, which is only returned when <code>binary = TRUE</code> .

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Friedman, J. H. (1991). "Multivariate adaptive regression splines." *Ann. Statist.* **19** 1–141.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.

Examples

```
data = friedman(100, 10, 1, FALSE)
```

mc.abc.vs

Variable selection with ABC Bayesian forest (using parallel computation)

Description

This function implements the variable selection approach proposed in Liu, Rockova and Wang (2021) with parallel computation. Rockova and Pas (2020) introduce a spike-and-forest prior which wraps the BART prior with a spike-and-slab prior on the model space. Due to intractable marginal likelihood, Liu, Rockova and Wang (2021) propose an approximate Bayesian computation (ABC) sampling method based on data-splitting to help sample from the model space with higher ABC acceptance rate.

Unlike the function `abc.vs()` which sequentially evaluates ABC models, this function evaluates the ABC models in parallel.

Usage

```

mc.abc.vs(
  x,
  y,
  nabc = 1000,
  tolerance = 0.1,
  threshold = 0.25,
  beta.params = c(1, 1),
  split.ratio = 0.5,
  probit = FALSE,
  true.idx = NULL,
  sparse = FALSE,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
  rm.const = TRUE,
  k = 2,
  power = 2,
  base = 0.95,
  split.prob = "polynomial",
  ntree = 10L,
  ndpost = 1,
  nskip = 200,
  keepevery = 1L,
  printevery = 100L,
  verbose = FALSE,
  mc.cores = 2L,
  nice = 19L,
  seed = 99L
)

```

Arguments

x	A matrix or a data frame of predictors values with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
y	A vector of response (continuous or binary) values.
nabc	The number of ABC samples, i.e., the number of subsets sampled from the model space.
tolerance	A number between 0 and 1; the nabc subsets are ranked by MSE in ascending order if the response variable is continuous (or by mean log loss (MLL) if the response variable is binary), and the top $\text{tolerance} \times 100\%$ of the subsets are accepted by ABC for selection.
threshold	A number between 0 and 1; within the ABC accepted subsets, predictors with MPVIP exceeding threshold are selected.

beta.params	A vector with two positive numbers; the spike-and-slab prior on the model space is assumed to be a beta-binomial prior, i.e., $\theta \text{Beta}(\text{beta.params}[1], \text{beta.params}[2])$ and each predictor is included into a model by $\text{Bernoulli}(\theta)$.
split.ratio	A number between 0 and 1; the data set (x, y) is split into a training set and a testing set according to the <code>split.ratio</code> .
probit	A Boolean argument indicating whether the response variable is binary or continuous; <code>probit=FALSE</code> (by default) means that the response variable is continuous.
true.idx	(Optional) A vector of indices of the true relevant predictors; if <code>true.idx</code> is provided, metrics including precision, recall and F1 score are returned.
sparse	A Boolean argument indicating whether to perform DART or BART.
xinfo	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
numcut	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in x . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of x . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x - 1)</code> cut-point values are used.
usequants	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
cont	A Boolean argument indicating whether to assume all predictors are continuous.
rm.const	A Boolean argument indicating whether to remove constant predictors.
k	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (y) is internally scaled to the range from $-.5$ to $.5$. The bigger k is, the more conservative the fitting will be.
power	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
base	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
split.prob	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
nree	The number of trees in the ensemble.
ndpost	The number of posterior samples returned.
nskip	The number of posterior samples burned in.
keepevery	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
printevery	As the MCMC runs, a message is printed every <code>printevery</code> iterations.

verbose	A Boolean argument indicating whether any messages are printed out.
mc.cores	The number of cores to employ in parallel.
nice	Set the job niceness. The default niceness is 19 and niceness goes from 0 (highest) to 19 (lowest).
seed	Seed required for reproducible MCMC.

Details

At each iteration of the algorithm, the data set is randomly split into a training set and a testing set according to a certain split ratio. The algorithm proceeds by sampling a subset from the spike-and-slab prior on the model space, fitting a BART model on the training set only with the predictors in the subset, and computing the root mean squared errors (RMSE) for the test set based on a posterior sample from the fitted BART model. Only those subsets that result in a low RMSE on the test set are kept for selection. ABC Bayesian forest selects predictors based on their marginal posterior variable inclusion probabilities (MPVIPs) which are estimated by computing the proportion of ABC accepted BART posterior samples that use the predictor at least one time. Given the MPVIPs, predictors with MPVIP exceeding a pre-specified threshold are selected. See Liu, Rockova and Wang (2021) or Section 2.2.4 in Luo and Daniels (2021) for details.

Value

The function `mc.abc.vs()` returns a list with the following components.

theta	The probability that a predictor is included into a model.
models	A matrix with <code>nabc</code> rows and <code>ncol(x)</code> columns; each row corresponds to a ABC model (or subset); if the (i, j) -th element is 1, it means that the j -th predictor is included in the i -th ABC model; if the (i, j) -th element is 0, it means that the j -th predictor is not included in the i -th ABC model.
actual.models	A matrix with <code>nabc</code> rows and <code>ncol(x)</code> columns; each row corresponds to a ABC BART posterior sample; if the (i, j) -th element is 1, it means that the j -th predictor is used as a split variable at least one time in the BART posterior sample of the i -th ABC model; if the (i, j) -th element is 0, it means that the j -th predictor is not used as a split variable in the BART posterior sample of the i -th ABC model.
model.errors	The vector of MSEs (or MLLs if the response variable is binary) for the <code>nabc</code> ABC models.
idx	The vector of indices (in terms of the row numbers of <code>models</code>) of the ABC accepted models which are the top <code>tolerance*100%</code> of the <code>nabc</code> ABC models when ranked by MSE or MLL in ascending order.
top.models	A matrix with <code>length(idx)</code> rows and <code>ncol(x)</code> columns, representing the ABC accepted models; <code>top.models=models[idx,]</code> .
top.actual.models	A matrix with <code>length(idx)</code> rows and <code>ncol(x)</code> columns, representing the ABC accepted BART posterior samples; <code>top.models=actual.models[idx,]</code> .
mip	The vector of marginal posterior variable inclusion probabilities.
best.model	The vector of predictors selected by ABC Bayesian forest.

precision	The precision score for the ABC Bayesian forest; only returned when true.idx is provided.
recall	The recall score for the ABC Bayesian forest; only returned when true.idx is provided.
f1	The F1 score for the ABC Bayesian forest; only returned when true.idx is provided.

Author(s)

Chuji Luo: <cjluc@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.

Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.

Liu, Yi, Veronika Rockova, and Yuexi Wang (2021). "Variable selection with ABC Bayesian forests." *J. R. Stat. Soc. Ser. B. Stat. Methodol.* 83.3, pp. 453–481.

Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.

Rockova Veronika and Stephanie van der Pas (2020). "Posterior concentration for Bayesian regression trees and forests." *Ann. Statist.* 48.4, pp. 2108–2131.

See Also

[abc.vs](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## parallel::mcparrallel/mccollect do not exist on windows
if(.Platform$OS.type=='unix') {
## test mc.abc.vs() function
  res = mc.abc.vs(data$X, data$Y, nabc=100, tolerance=0.1, threshold=0.25,
  beta.params=c(1.0, 1.0), split.ratio=0.5, probit=FALSE, true.idx=c(1,2,6:8),
  ntree=10, ndpost=1, nskip=200, mc.cores=2)
}
```

`mc.backward.vs`*Backward selection with two filters (using parallel computation)*

Description

This function implements the backward variable selection approach for BART (see Algorithm 2 in Luo and Daniels (2021) for details). Parallel computation is used within each step of the backward selection approach.

Usage

```
mc.backward.vs(  
  x,  
  y,  
  split.ratio = 0.8,  
  probit = FALSE,  
  true.idx = NULL,  
  xinfo = matrix(0, 0, 0),  
  numcut = 100L,  
  usequants = FALSE,  
  cont = FALSE,  
  rm.const = TRUE,  
  k = 2,  
  power = 2,  
  base = 0.95,  
  split.prob = "polynomial",  
  ntree = 50L,  
  ndpost = 1000,  
  nskip = 1000,  
  keepevery = 1L,  
  printevery = 100L,  
  verbose = FALSE,  
  mc.cores = 2L,  
  nice = 19L,  
  seed = 99L  
)
```

Arguments

<code>x</code>	A matrix or a data frame of predictors values with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y</code>	A vector of response (continuous or binary) values.
<code>split.ratio</code>	A number between 0 and 1; the data set (x, y) is split into a training set and a testing set according to the <code>split.ratio</code> .

probit	A Boolean argument indicating whether the response variable is binary or continuous; <code>probit=FALSE</code> (by default) means that the response variable is continuous.
true.idx	(Optional) A vector of indices of the true relevant predictors; if provided, metrics including precision, recall and F1 score are returned.
xinfo	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
numcut	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in <code>x</code> . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x - 1)</code> cut-point values are used.
usequants	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
cont	A Boolean argument indicating whether to assume all predictors are continuous.
rm.const	A Boolean argument indicating whether to remove constant predictors.
k	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y</code>) is internally scaled to the range from $-.5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
power	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
base	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
split.prob	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
ntree	The number of trees in the ensemble.
ndpost	The number of posterior samples returned.
nskip	The number of posterior samples burned in.
keepevery	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
printevery	As the MCMC runs, a message is printed every <code>printevery</code> iterations.
verbose	A Boolean argument indicating whether any messages are printed out.
mc.cores	The number of cores to employ in parallel.
nice	Set the job niceness. The default niceness is 19 and niceness goes from 0 (highest) to 19 (lowest).
seed	Seed required for reproducible MCMC.

Details

The backward selection starts with the full model with all the predictors, followed by comparing the deletion of each predictor using mean squared error (MSE) if the response variable is continuous (or mean log loss (MLL) if the response variable is binary) and then deleting the predictor whose loss gives the smallest MSE (or MLL). This process is repeated until there is only one predictor in the model and ultimately returns `ncol{x}` "winner" models with different model sizes ranging from 1 to `ncol{x}`.

Given the `ncol{x}` "winner" models, the one with the largest expected log pointwise predictive density based on leave-one-out (LOO) cross validation is the best model. See Section 3.3 in Luo and Daniels (2021) for details.

If `true.idx` is provided, the precision, recall and F1 scores are returned.

Value

The function `mc.backward.vs()` returns a list with the following components.

<code>best.model.names</code>	The vector of column names of the predictors selected by the backward selection approach.
<code>best.model.cols</code>	The vector of column indices of the predictors selected by the backward selection approach.
<code>best.model.order</code>	The step where the best model is located.
<code>models</code>	The list of winner models from each step of the backward selection procedure; length equals <code>ncol{x}</code> .
<code>model.errors</code>	The vector of MSEs (or MLLs if the response variable is binary) for the <code>ncol{x}</code> winner models.
<code>elpd.loos</code>	The vector of LOO scores for the <code>ncol{x}</code> winner models.
<code>all.models</code>	The list of all the evaluated models.
<code>all.model.errors</code>	The vector of MSEs (or MLLs if the response variable is binary) for all the evaluated models.
<code>precision</code>	The precision score for the backward selection approach; only returned if <code>true.idx</code> is provided.
<code>recall</code>	The recall score for the backward selection approach; only returned if <code>true.idx</code> is provided.
<code>f1</code>	The F1 score for the backward selection approach; only returned if <code>true.idx</code> is provided.
<code>all.models.idx</code>	The vector of Boolean arguments indicating whether the corresponding model in <code>all.models</code> is acceptable or not; a model containing all the relevant predictors is an acceptable model; only returned if <code>true.idx</code> is provided.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Vehtari, Aki, Andrew Gelman, and Jonah Gabry (2017). "Erratum to: Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC." *Stat. Comput.* **27.5**, p. 1433.

See Also

[permute.vs](#), [medianInclusion.vs](#) and [abc.vs](#).

Examples

```
## simulate data (Scenario C.C.1. in Luo and Daniels (2021))
set.seed(123)
data = friedman(100, 5, 1, FALSE)
## parallel::mcpParallel/mccollect do not exist on windows
if(.Platform$OS.type=='unix') {
## test mc.backward.vs() function
  res = mc.backward.vs(data$X, data$Y, split.ratio=0.8, probit=FALSE,
    true.idx=c(1:5), ntree=10, ndpost=100, nskip=100, mc.cores=2)
}
```

mc.cores.openmp

Detecting OpenMP

Description

This function is inherited from the CRAN R package 'BART' and was designed for OpenMP. For example, the `pwbart` function can use OpenMP or the 'parallel' R package for multi-threading. On UNIX/Unix-like systems, OpenMP, if available, is discovered at install time. However, we know of no GPL licensed code available to detect OpenMP on Windows (for Artistic licensed OpenMP detection code on Windows, see the Bioconductor R package 'rGADEM'). To determine whether OpenMP is available at run time, we provide the function documented here.

Usage

```
mc.cores.openmp()
```

Value

This function returns 0 when OpenMP is not available; otherwise, an integer greater than 0 is returned when OpenMP is available (1 is returned unless you are running in a multi-threaded process)

Author(s)

Chuji Luo: <cjl1uo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.

Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.

Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.

Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.

Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[pwbart](#).

Examples

```
mc.cores.openmp()
```

```
mc.pbart
```

Probit BART for binary responses with parallel computation

Description

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a binary response y and a p -dimensional vector of predictors $x = (x_1, \dots, x_p)'$, probit BART models y and x using

$$P(Y = 1|x) = \Phi[f(x)],$$

where Φ is the CDF of the standard normal distribution and f is a sum of Bayesian regression trees function.

The function `mc.pbart()` is inherited from the CRAN R package 'BART' and is a variant of the function `pbart()` with parallel computation.

Usage

```

mc.pbart(
  x.train,
  y.train,
  x.test = matrix(0, 0L, 0L),
  sparse = FALSE,
  theta = 0,
  omega = 1,
  a = 0.5,
  b = 1,
  augment = FALSE,
  rho = NULL,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
  rm.const = TRUE,
  k = 2,
  power = 2,
  base = 0.95,
  split.prob = "polynomial",
  binaryOffset = NULL,
  ntree = 50L,
  ndpost = 1000L,
  nskip = 100L,
  keepevery = 1L,
  printevery = 100L,
  keeptrainfits = TRUE,
  transposed = FALSE,
  verbose = FALSE,
  mc.cores = 2L,
  nice = 19L,
  seed = 99L
)

```

Arguments

<code>x.train</code>	A matrix or a data frame of predictors values (for training) with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y.train</code>	A vector of continuous response values for training.
<code>x.test</code>	A matrix or a data frame of predictors values for testing, which has the same structure as <code>x.train</code> .
<code>sparse</code>	A Boolean argument indicating whether to replace the discrete uniform distribution for selecting a split variable with a categorical distribution whose event probabilities follow a Dirichlet distribution (see Linero (2018) for details).

theta	Set theta parameter; zero means random.
omega	Set omega parameter; zero means random.
a	A sparse parameter of $Beta(a, b)$ hyper-prior where $0.5 \leq a \leq 1$; a lower value induces more sparsity.
b	A sparse parameter of $Beta(a, b)$ hyper-prior; typically, $b = 1$.
augment	A Boolean argument indicating whether data augmentation is performed in the variable selection procedure of Linero (2018).
rho	A sparse parameter; typically $\rho = p$ where p is the number of predictors.
xinfo	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
numcut	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x.train)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in <code>x.train</code> . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x.train</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x.train - 1)</code> cut-point values are used.
usequants	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
cont	A Boolean argument indicating whether to assume all predictors are continuous.
rm.const	A Boolean argument indicating whether to remove constant predictors.
k	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y.train</code>) is internally scaled to the range from $-.5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
power	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
base	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
split.prob	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
binaryOffset	The binary offset term in the probit BART model, i.e., $P(Y = 1 x) = \Phi[f(x) + \text{binaryOffset}]$.
nree	The number of trees in the ensemble.
ndpost	The number of posterior samples returned.
nskip	The number of posterior samples burned in.
keepevery	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
printevery	As the MCMC runs, a message is printed every <code>printevery</code> iterations.

keeptrainfits	A Boolean argument indicating whether to keep <code>yhat.train</code> or not.
transposed	A Boolean argument indicating whether the matrices <code>x.train</code> and <code>x.test</code> are transposed.
verbose	A Boolean argument indicating whether any messages are printed out.
mc.cores	The number of cores to employ in parallel.
nice	Set the job niceness. The default niceness is 19 and niceness goes from 0 (highest) to 19 (lowest).
seed	Seed required for reproducible MCMC.

Details

This function is inherited from `BART::mc.pbart()` and is a variant of the function `pbart()` with parallel computation. While the original features of `BART::pbart()` are preserved, two modifications are made.

The first modification is to provide two types of split probability for BART. One split probability is proposed in Chipman et al. (2010) and defined as

$$p(d) = \gamma * (1 + d)^{-\beta},$$

where d is the depth of the node, $\gamma \in (0, 1)$ and $\beta \in (0, \infty)$. The other split probability is proposed by Rockova and Saha (2019) and defined as

$$p(d) = \gamma^d,$$

where $\gamma \in (1/n, 1/2)$. BART with the second split probability is proved to achieve the optimal posterior contraction.

The second modification is to provide five types of variable importance measures (`vip`, `within.type.vip`, `pvip`, `varprob.mean` and `mi`) in the return object, for the sake of the existence of mixed-type predictors.

Value

The function `pbart()` returns an object of type `pbart` which essentially is a list consisting of the following components.

<code>yhat.train</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.train)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a training data point. The (i, j) -th element of the matrix is $f^*(x) + \text{binaryOffset}$ for the i -th kept draw of f and the j -th training data point. Burn-in posterior samples are dropped.
<code>prob.train</code>	A matrix with the same structure as <code>yhat.train</code> and the (i, j) -th element of <code>prob.train</code> is $\Phi[f^*(x) + \text{binaryOffset}]$ for the i -th kept draw of f and the j -th training data point.
<code>prob.train.mean</code>	A vector which is <code>colMeans(prob.train)</code> .
<code>yhat.test</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.test)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a test data point. The (i, j) -th element of the matrix is $f^*(x) + \text{binaryOffset}$ for the i -th kept draw of f and the j -th test data point. Burn-in posterior samples are dropped.

prob.test	A matrix with the same structure as yhat.test and the (i, j) -th element of prob.test is $\Phi[f * (x) + \text{binaryOffset}]$ for the i -th kept draw of f and the j -th test data point.
prob.test.mean	A vector which is colMeans(prob.test).
varcount	A matrix with ndpost rows and ncol(x.train) columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the number of times that the j -th predictor is used as a split variable in the i -th posterior sample.
varprob	A matrix with ndpost rows and ncol(x.train) columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the split probability of the j -th predictor in the i -th posterior sample. Only useful when DART is fit, i.e., sparse=TRUE.
treedraws	A list containing the posterior samples of the ensembles (trees structures, split variables and split values); Can be used for prediction.
proc.time	The process time of running the function wbart().
vip	A vector of variable inclusion proportions (VIP) proposed in Chipman et al. (2010).
within.type.vip	A vector of within-type VIPs proposed in Luo and Daniels (2021).
pvip	A vector of marginal posterior variable inclusion probabilities (PVIP) proposed in Linero (2018); Only useful when DART is fit, i.e., sparse=TRUE.
varprob.mean	A vector of posterior split probabilities (PSP) proposed in Linero (2018); Only useful when DART is fit, i.e., sparse=TRUE.
mr.vecs	A list of $ncol(x.train)$ sub-lists with each corresponding to a predictor; Each sub-list contains ndpost vectors with each vector containing the (birth) Metropolis ratios for splits using the predictor as the split variable in that posterior sample.
mr.mean	A matrix with ndpost rows and ncol(x.train) columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the average Metropolis acceptance ratio per splitting rule using the j -th predictor in the i -th posterior sample.
mi	A vector of Metropolis importance (MI) proposed in Luo and Daniels (2021).
rm.const	A vector of indicators for the predictors (after dummification) used in BART; when the indicator is negative, it refers to remove that predictor.
binaryOffset	The binary offset term used in BART.
ndpost	The number of posterior samples returned.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[pbart](#).

Examples

```
## simulate data (Scenario B.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, TRUE)
## parallel::mcp/parallel/mccollect do not exist on windows
if(.Platform$OS.type=='unix') {
## test mc.pbart() function
  res = mc.pbart(data$X, data$Y, ntree=10, nskip=100, ndpost=100, mc.cores=2)
}
```

mc.permute.vs

Permutation-based variable selection approach with parallel computation

Description

This function implements the permutation-based variable selection approach for BART (see Algorithm 1 in Luo and Daniels (2021) for details) with parallel computation used in computing the null variable importance scores. Three types of variable importance measures are considered: BART variable inclusion proportions (VIP), BART within-type variable inclusion proportions (within-type VIP) and BART Metropolis Importance (MI).

The permutation-based variable selection approach using BART VIP as the variable importance measure is proposed by Bleich et al. (2014). BART within-type VIP and BART MI are proposed by Luo and Daniels (2021), for the sake of the existence of mixed-type predictors and the goal of allowing more relevant predictors into the model.

Usage

```

mc.permute.vs(
  x.train,
  y.train,
  probit = FALSE,
  npermute = 100L,
  nreps = 10L,
  alpha = 0.05,
  true.idx = NULL,
  plot = TRUE,
  n.var.plot = Inf,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
  rm.const = TRUE,
  k = 2,
  power = 2,
  base = 0.95,
  split.prob = "polynomial",
  ntree = 20L,
  ndpost = 1000,
  nskip = 1000,
  keepevery = 1L,
  verbose = FALSE,
  mc.cores = 2L,
  nice = 19L,
  seed = 99L
)

```

Arguments

<code>x.train</code>	A matrix or a data frame of predictors values with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y.train</code>	A vector of response (continuous or binary) values.
<code>probit</code>	A Boolean argument indicating whether the response variable is binary or continuous; <code>probit=FALSE</code> (by default) means that the response variable is continuous.
<code>npermute</code>	The number of permutations for estimating the null distributions of the variable importance scores.
<code>nreps</code>	The number of replications for obtaining the averaged (or median) variable importance scores based on the original data set.
<code>alpha</code>	A number between 0 and 1; a predictor is selected if its averaged (or median) variable importance score exceeds the $1 - \alpha$ quantile of the corresponding null distribution.

<code>true.idx</code>	(Optional) A vector of indices of the true relevant predictors; if provided, metrics including precision, recall and F1 score will be returned.
<code>plot</code>	(Optional) A Boolean argument indicating whether plots are returned or not.
<code>n.var.plot</code>	The number of variables to be plotted.
<code>xinfo</code>	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
<code>numcut</code>	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x.train)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in <code>x.train</code> . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x.train</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x.train - 1)</code> cut-point values are used.
<code>usequants</code>	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
<code>cont</code>	A Boolean argument indicating whether to assume all predictors are continuous.
<code>rm.const</code>	A Boolean argument indicating whether to remove constant predictors.
<code>k</code>	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y.train</code>) is internally scaled to the range from $-.5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
<code>power</code>	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
<code>base</code>	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
<code>split.prob</code>	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
<code>ntree</code>	The number of trees in the ensemble.
<code>ndpost</code>	The number of posterior samples returned.
<code>nskip</code>	The number of posterior samples burned in.
<code>keepevery</code>	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
<code>verbose</code>	A Boolean argument indicating whether any messages are printed out.
<code>mc.cores</code>	The number of cores to employ in parallel.
<code>nice</code>	Set the job niceness. The default niceness is 19 and niceness goes from 0 (highest) to 19 (lowest).
<code>seed</code>	Seed required for reproducible MCMC.

Details

The detailed algorithm can be found in Algorithm 1 in Luo and Daniels (2021). The permutation-based variable selection approach using within-type VIP as the variable importance measure is only used when the predictors are of mixed-type; otherwise, it is the same as the one using VIP as the variable importance measure.

If `true.idx` is provided, the precision, recall and F1 scores will be returned for the three (or two if the predictors are of the same type) methods.

If `plot=TRUE`, three (or two if the predictors are of the same type) plots showing which predictors are selected are generated.

Value

The function `mc.permute.vs()` returns three (or two if the predictors are of the same type) plots if `plot=TRUE` and a list with the following components.

<code>vip.imp.cols</code>	The vector of column indices of the predictors selected by the approach using VIP as the variable importance score.
<code>vip.imp.names</code>	The vector of column names of the predictors selected by the approach using VIP as the variable importance score.
<code>avg.vip</code>	The vector (<code>length=ncol(x.train)</code>) of the averaged VIPs based on the original data set; <code>avg.vip=colMeans(avg.vip.mtx)</code> .
<code>avg.vip.mtx</code>	A matrix of VIPs based on the original data set, with each row corresponding to a repetition and each column corresponding to a predictor.
<code>permute.vips</code>	A matrix of VIPs based on the null data sets, with each row corresponding to a permutation (null data set) and each column corresponding to a predictor.
<code>within.type.vip.cols</code>	The vector of column indices of the predictors selected by the approach using within-type VIP as the variable importance score.
<code>within.type.vip.names</code>	The vector of column names of the predictors selected by the approach using within-type VIP as the variable importance score.
<code>avg.within.type.vip</code>	The vector (<code>length=ncol(x.train)</code>) of the averaged within-type VIPs based on the original data set; <code>avg.within.type.vip=colMeans(avg.within.type.vip.mtx)</code> .
<code>avg.within.type.vip.mtx</code>	A matrix of within-type VIPs based on the original data set, with each row corresponding to a repetition and each column corresponding to a predictor.
<code>permute.within.type.vips</code>	A matrix of within VIPs based on the null data sets, with each row corresponding to a permutation (null data set) and each column corresponding to a predictor.
<code>mi.imp.cols</code>	The vector of column indices of the predictors selected by the approach using MI as the variable importance score.
<code>mi.imp.names</code>	The vector of column names of the predictors selected by the approach using MI as the variable importance score.
<code>median.mi</code>	The vector (<code>length=ncol(x.train)</code>) of the median MIs based on the original data set; <code>median.mi=colMeans(median.mi.mtx)</code> .

median.mi.mtx	A matrix of MIs based on the original data set, with each row corresponding to a repetition and each column corresponding to a predictor.
permute.mis	A matrix of MIs based on the null data sets, with each row corresponding to a permutation (null data set) and each column corresponding to a predictor.
true.idx	A vector of indices of the true relevant predictors; only returned if true.idx is provided as inputs.
vip.precision	The precision score for the approach using VIP as the variable importance score; only returned if true.idx is provided.
vip.recall	The recall score for the approach using VIP as the variable importance score; only returned if true.idx is provided.
vip.f1	The F1 score for the approach using VIP as the variable importance score; only returned if true.idx is provided.
wt.vip.precision	The precision score for the approach using within-VIP as the variable importance score; only returned when the predictors are of the same type and true.idx is provided.
wt.vip.recall	The recall score for the approach using within-VIP as the variable importance score; only returned when the predictors are of the same type and true.idx is provided.
wt.vip.f1	The F1 score for the approach using within-VIP as the variable importance score; only returned when the predictors are of the same type and true.idx is provided.
mi.precision	The precision score for the approach using MI as the variable importance score; only returned if true.idx is provided.
mi.recall	The recall score for the approach using MI as the variable importance score; only returned if true.idx is provided.
mi.f1	The F1 score for the approach using MI as the variable importance score; only returned if true.idx is provided.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Bleich, Justin et al. (2014). "Variable selection for BART: an application to gene regulation." *Ann. Appl. Stat.* 8.3, pp 1750–1781.
- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* 4 266–298.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.

See Also

[permute.vs](#), [medianInclusion.vs](#), [mc.backward.vs](#) and [abc.vs](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## parallel::mcparrallel/mccollect do not exist on windows
if(.Platform$OS.type=='unix') {
## test mc.permute.vs() function
  res = mc.permute.vs(data$X, data$Y, probit=FALSE, npermute=100, nreps=10, alpha=0.05,
  true.idx=c(1, 2, 6:8), plot=FALSE, ntree=10, ndpost=100, nskip=100, mc.cores=2)
}
```

mc.pwbart

Predicting new observations based on a previously fitted BART model with parallel computation

Description

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a continuous response y and a p -dimensional vector of predictors $x = (x_1, \dots, x_p)'$, BART models y and x using

$$y = f(x) + \epsilon,$$

where f is a sum of Bayesian regression trees function and $\epsilon \sim N(0, \sigma^2)$.

For a binary response y , probit BART models y and x using

$$P(Y = 1|x) = \Phi[f(x)],$$

where Φ is the CDF of the standard normal distribution and f is a sum of Bayesian regression trees function.

The function `mc.pwbart()` is inherited from the CRAN R package 'BART'.

Usage

```
mc.pwbart(
  x.test,
  treedraws,
  rm.const,
  mu = 0,
  mc.cores = 2L,
  transposed = FALSE,
  dodraws = TRUE,
  nice = 19L
)
```

Arguments

x.test	A matrix or a data frame of predictors values for prediction with each row corresponding to an observation and each column corresponding to a predictor.
treedraws	A list which is the \$treedraws returned from the function wbart() or pbart().
rm.const	A vector which is the \$rm.const returned from the function wbart() or pbart().
mu	Mean to add on to y prediction.
mc.cores	The number of threads to utilize.
transposed	A Boolean argument indicating whether the matrix x.test is transposed. When running pwbart() or mc.pwbart() in parallel, it is more memory-efficient to transpose x.test prior to calling the internal versions of these functions.
dodraws	A Boolean argument indicating whether to return the draws themselves (the default), or whether to return the mean of the draws as specified by dodraws=FALSE.
nice	Set the job niceness. The default niceness is 19 and niceness goes from 0 (highest) to 19 (lowest).

Value

Returns the predictions for x.test. If dodraws=TRUE, return a matrix of prediction with each row corresponding to a draw and each column corresponding to a new observation; if dodraws=FALSE, return a vector of predictions which are the mean of the draws.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[wbart](#), [pbart](#) and [mc.pwbart](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## run wbart() function
res = wbart(data$X, data$Y, ntree=10, nskip=100, ndpost=100)
## parallel::mcpParallel/mcCollect do not exist on windows
if(.Platform$OS.type=='unix') {
## test pwbart() function
  x.test = mixone(5, 10, 1, FALSE)$X
  pred = mc.pwbart(x.test, res$treedraws, res$rm.const, mu=mean(data$Y), mc.cores=2)
}
```

mc.wbart

BART for continuous responses with parallel computation

Description

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a continuous response y and a p -dimensional vector of predictors $x = (x_1, \dots, x_p)'$, BART models y and x using

$$y = f(x) + \epsilon,$$

where f is a sum of Bayesian regression trees function and $\epsilon \sim N(0, \sigma^2)$.

The function `mc.wbart()` is inherited from the CRAN R package 'BART' and is a variant of the function `wbart()` with parallel computation.

Usage

```
mc.wbart(
  x.train,
  y.train,
  x.test = matrix(0, 0, 0),
  sparse = FALSE,
  theta = 0,
  omega = 1,
  a = 0.5,
  b = 1,
  augment = FALSE,
  rho = NULL,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
  rm.const = TRUE,
  power = 2,
```



```

base = 0.95,
split.prob = "polynomial",
k = 2,
sigmaf = NA,
sigest = NA,
sigdf = 3,
sigquant = 0.9,
lambda = NA,
fmean = mean(y.train),
w = rep(1, length(y.train)),
ntree = 200L,
ndpost = 1000L,
nskip = 100L,
keepevery = 1L,
printevery = 100L,
keeptrainfits = TRUE,
transposed = FALSE,
verbose = FALSE,
mc.cores = 2L,
nice = 19L,
seed = 99L
)

```

Arguments

<code>x.train</code>	A matrix or a data frame of predictors values (for training) with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y.train</code>	A vector of continuous response values for training.
<code>x.test</code>	A matrix or a data frame of predictors values for testing, which has the same structure as <code>x.train</code> .
<code>sparse</code>	A Boolean argument indicating whether to replace the discrete uniform distribution for selecting a split variable with a categorical distribution whose event probabilities follow a Dirichlet distribution (see Linero (2018) for details).
<code>theta</code>	Set theta parameter; zero means random.
<code>omega</code>	Set omega parameter; zero means random.
<code>a</code>	A sparse parameter of $Beta(a, b)$ hyper-prior where $0.5 \leq a \leq 1$; a lower value induces more sparsity.
<code>b</code>	A sparse parameter of $Beta(a, b)$ hyper-prior; typically, $b = 1$.
<code>augment</code>	A Boolean argument indicating whether data augmentation is performed in the variable selection procedure of Linero (2018).
<code>rho</code>	A sparse parameter; typically $\rho = p$ where p is the number of predictors.
<code>xinfo</code>	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.

numcut	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x.train)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in <code>x.train</code> . If <code>usequants=FALSE</code> , numcut equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x.train</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x.train - 1)</code> cut-point values are used.
usequants	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
cont	A Boolean argument indicating whether to assume all predictors are continuous.
rm.const	A Boolean argument indicating whether to remove constant predictors.
power	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
base	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
split.prob	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
k	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y.train</code>) is internally scaled to the range from $-.5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
sigmaf	The standard deviation of <code>f</code> .
sigest	A rough estimate of the error standard deviation, the square of which follows an inverse chi-squared prior. If <code>sigest=NA</code> , the rough estimate will be the usual least square estimator; Otherwise, the supplied value will be used.
sigdf	The degrees of freedom for the error variance prior.
sigquant	The quantile of the error variance prior, where <code>sigest</code> is placed. The closer the quantile is to 1, the more aggressive the fit will be.
lambda	The scale parameter of the error variance prior.
fmean	BART operates on <code>y.train</code> centered by <code>fmean</code> .
w	A vector of weights which multiply the standard deviation.
ntree	The number of trees in the ensemble.
ndpost	The number of posterior samples returned.
nskip	The number of posterior samples burned in.
keepevery	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
printevery	As the MCMC runs, a message is printed every <code>printevery</code> iterations.
keeptrainfits	A Boolean argument indicating whether to keep <code>yhat.train</code> or not.
transposed	A Boolean argument indicating whether the matrices <code>x.train</code> and <code>x.test</code> are transposed.

verbose	A Boolean argument indicating whether any messages are printed out.
mc.cores	The number of cores to employ in parallel.
nice	Set the job niceness. The default niceness is 19 and niceness goes from 0 (highest) to 19 (lowest).
seed	Seed required for reproducible MCMC.

Details

This function is inherited from `BART::mc.wbart()` and is a variant of the function `wbart()` with parallel computation.

While the original features of `BART::wbart()` are preserved, two modifications are made.

The first modification is to provide two types of split probability for BART. One split probability is proposed in Chipman et al. (2010) and defined as

$$p(d) = \gamma * (1 + d)^{-\beta},$$

where d is the depth of the node, $\gamma \in (0, 1)$ and $\beta \in (0, \infty)$. The other split probability is proposed by Rockova and Saha (2019) and defined as

$$p(d) = \gamma^d,$$

where $\gamma \in (1/n, 1/2)$. BART with the second split probability is proved to achieve the optimal posterior contraction.

The second modification is to provide five types of variable importance measures (`vip`, `within.type.vip`, `pvip`, `varprob.mean` and `mi`) in the return object, for the sake of the existence of mixed-type predictors.

Value

The function `mc.wbart()` returns an object of type `wbart` which essentially is a list consisting of the following components.

<code>sigma</code>	A vector with <code>nskip+ndpost*keepevery</code> posterior samples of σ .
<code>yhat.train.mean</code>	<code>colMeans(yhat.train)</code> .
<code>yhat.train</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.train)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a training data point. The (i, j) -th element of the matrix is $f^*(x)$ for the i -th kept draw of f and the j -th training data point. Burn-in posterior samples are dropped.
<code>yhat.test.mean</code>	<code>colMeans(yhat.test)</code> .
<code>yhat.test</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.test)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a test data point. The (i, j) -th element of the matrix is $f^*(x)$ for the i -th kept draw of f and the j -th test data point. Burn-in posterior samples are dropped.
<code>varcount</code>	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the number of times that the j -th predictor is used as a split variable in the i -th posterior sample.

varprob	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the split probability of the j -th predictor in the i -th posterior sample. Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
treedraws	A list containing the posterior samples of the ensembles (trees structures, split variables and split values); Can be used for prediction.
proc.time	The process time of running the function <code>wbart()</code> .
mu	BART operates on <code>y.train</code> centered by <code>fmean</code> .
mr.vecs	A list of <code>ncol(x.train)</code> sub-lists with each corresponding to a predictor; Each sub-list contains <code>ndpost</code> vectors with each vector containing the (birth) Metropolis ratios for splits using the predictor as the split variable in that posterior sample.
vip	A vector of variable inclusion proportions (VIP) proposed in Chipman et al. (2010).
within.type.vip	A vector of within-type VIPs proposed in Luo and Daniels (2021).
pvip	A vector of marginal posterior variable inclusion probabilities (PVIP) proposed in Linero (2018); Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
varprob.mean	A vector of posterior split probabilities (PSP) proposed in Linero (2018); Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
mr.mean	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the average Metropolis acceptance ratio per splitting rule using the j -th predictor in the i -th posterior sample.
mi	A vector of Metropolis importance (MI) proposed in Luo and Daniels (2021).
rm.const	A vector of indicators for the predictors (after dummification) used in BART; when the indicator is negative, it refers to remove that predictor.
ndpost	The number of posterior samples returned.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also[wbart.](#)**Examples**

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## parallel::mcpipeline/mccollect do not exist on windows
if(.Platform$OS.type=='unix') {
  ## test mc.wbart() function
  res = mc.wbart(data$X, data$Y, ntree=100, nskip=100, ndpost=100, mc.cores=2)
}
```

medianInclusion.vs *Variable selection with DART*

Description

This function implements the variable selection approach proposed in Linero (2018). Linero (2018) proposes DART, a variant of BART, which replaces the discrete uniform distribution for selecting a split variable with a categorical distribution of which the event probabilities follow a Dirichlet distribution. DART estimates the marginal posterior variable inclusion probability (MPVIP) for a predictor by the proportion of the posterior samples of the trees structures where the predictor is used as a split variable at least once, and selects predictors with MPVIP at least 0.5, yielding a median probability model.

Usage

```
medianInclusion.vs(
  x.train,
  y.train,
  probit = FALSE,
  vip.selection = TRUE,
  true.idx = NULL,
  plot = FALSE,
  num.var.plot = Inf,
  theta = 0,
  omega = 1,
  a = 0.5,
  b = 1,
  augment = FALSE,
  rho = NULL,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
```

```

rm.const = TRUE,
power = 2,
base = 0.95,
split.prob = "polynomial",
k = 2,
ntree = 20L,
ndpost = 1000L,
nskip = 1000L,
keepevery = 1L,
printevery = 100L,
verbose = FALSE
)

```

Arguments

<code>x.train</code>	A matrix or a data frame of predictors values with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y.train</code>	A vector of response (continuous or binary) values.
<code>probit</code>	A Boolean argument indicating whether the response variable is binary or continuous; <code>probit=FALSE</code> (by default) means that the response variable is continuous.
<code>vip.selection</code>	A Boolean argument indicating whether to select predictors using BART VIPs.
<code>true.idx</code>	(Optional) A vector of indices of the true relevant predictors; if provided, metrics including precision, recall and F1 score are returned.
<code>plot</code>	(Optional) A Boolean argument indicating whether plots are returned or not.
<code>num.var.plot</code>	The number of variables to be plotted.
<code>theta</code>	Set theta parameter; zero means random.
<code>omega</code>	Set omega parameter; zero means random.
<code>a</code>	A sparse parameter of $Beta(a, b)$ hyper-prior where $0.5 \leq a \leq 1$; a lower value induces more sparsity.
<code>b</code>	A sparse parameter of $Beta(a, b)$ hyper-prior; typically, $b = 1$.
<code>augment</code>	A Boolean argument indicating whether data augmentation is performed in the variable selection procedure of Linero (2018).
<code>rho</code>	A sparse parameter; typically $\rho = p$ where p is the number of predictors.
<code>xinfo</code>	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
<code>numcut</code>	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x.train)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in <code>x.train</code> . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x.train</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x.train - 1)</code> cut-point values are used.

usequants	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
cont	A Boolean argument indicating whether to assume all predictors are continuous.
rm.const	A Boolean argument indicating whether to remove constant predictors.
power	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
base	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is base^d .
split.prob	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
k	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y.train</code>) is internally scaled to the range from $-.5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
nree	The number of trees in the ensemble.
ndpost	The number of posterior samples returned.
nskip	The number of posterior samples burned in.
keepevery	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
printevery	As the MCMC runs, a message is printed every <code>printevery</code> iterations.
verbose	A Boolean argument indicating whether any messages are printed out.

Details

See Linero (2018) or Section 2.2.3 in Luo and Daniels (2021) for details.

If `vip.selection=TRUE`, this function also does variable selection by selecting variables whose BART VIP exceeds $1/n_{\text{col}}\{x.\text{train}\}$.

If `true.idx` is provided, the precision, recall and F1 scores are returned.

If `plot=TRUE`, plots showing which predictors are selected are generated.

Value

The function `medianInclusion.vs()` returns two (or one if `vip.selection=FALSE`) plots if `plot=TRUE` and a list with the following components.

<code>dart.pvip</code>	The vector of DART MPVIPs.
<code>dart.pvip.imp.names</code>	The vector of column names of the predictors with DART MPVIP at least 0.5.
<code>dart.pvip.imp.cols</code>	The vector of column indices of the predictors with DART MPVIP at least 0.5.
<code>dart.precision</code>	The precision score for the DART approach; only returned if <code>true.idx</code> is provided.
<code>dart.recall</code>	The recall score for the DART approach; only returned if <code>true.idx</code> is provided.

<code>dart.f1</code>	The F1 score for the DART approach; only returned if <code>true.idx</code> is provided.
<code>bart.vip</code>	The vector of BART VIPs; only returned if <code>vip.selection=TRUE</code> .
<code>bart.vip.imp.names</code>	The vector of column names of the predictors with BART VIP exceeding $1/\text{ncol}\{x.\text{train}\}$; only returned if <code>vip.selection=TRUE</code> .
<code>bart.vip.imp.cols</code>	The vector of column indices of the predictors with BART VIP exceeding $1/\text{ncol}\{x.\text{train}\}$; only returned if <code>vip.selection=TRUE</code> .
<code>bart.precision</code>	The precision score for the BART approach; only returned if <code>vip.selection=TRUE</code> and <code>true.idx</code> is provided.
<code>bart.recall</code>	The recall score for the BART approach; only returned if <code>vip.selection=TRUE</code> and <code>true.idx</code> is provided.
<code>bart.f1</code>	The F1 score for the BART approach; only returned if <code>vip.selection=TRUE</code> and <code>true.idx</code> is provided.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.

See Also

[permute.vs](#), [mc.backward.vs](#) and [abc.vs](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## test medianInclusion.vs() function
res = medianInclusion.vs(data$X, data$Y, probit=FALSE, vip.selection=TRUE,
true.idx=c(1, 2, 6:8), plot=FALSE, ntree=10, ndpost=100, nskip=100, verbose=FALSE)
```

 mixone

 Generate data with independent and mixed-type predictors

Description

Generate data including responses and predictors values, of which predictors are independent and of mixed types.

Usage

```
mixone(n, p, sigma, binary)
```

Arguments

n	The number of observations.
p	The number of predictors.
sigma	The error variance.
binary	A boolean argument: <code>binary = TRUE</code> indicates that binary responses are generated and <code>binary = FALSE</code> indicates that continuous responses are generated.

Details

Sample the predictors $x_1, \dots, x_{\text{ceiling}(p/2)}$ from Bernoulli(0.5) independently and $x_{\text{ceiling}(p/2)+1}, \dots, x_p$ from Uniform(0, 1) independently. If `binary = FALSE`, sample the continuous response y from $\text{Normal}(f_0(x), \sigma^2)$, where

$$f_0(x) = 10\sin(\pi x_{\text{ceiling}(p/2)+1} * x_{\text{ceiling}(p/2)+2}) + 20(x_{\text{ceiling}(p/2)+3} - 0.5)^2 + 10x_1 + 5x_2.$$

If `binary = TRUE`, sample the binary response y from $\text{Bernoulli}(\Phi(f_0(x)))$ where f_0 is defined above and Φ is the cumulative density function of the standard normal distribution.

Value

Return a list with the following components.

X	An n by p data frame representing predictors values, with each row corresponding an observation.
Y	A vector of length n representing response values.
f0	A vector of length n representing the values of $f_0(x)$.
sigma	The error variance which is only returned when <code>binary = FALSE</code> .
prob	A vector of length n representing the values of $\Phi(f_0(x))$, which is only returned when <code>binary = TRUE</code> .

Author(s)

Chuji Luo: <cjluo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.

Examples

```
data = mixone(100, 10, 1, FALSE)
```

mixtwo	<i>Generate data with correlated and mixed-type predictors</i>
--------	--

Description

Generate data including responses and predictors values, of which predictors are correlated and of mixed types.

Usage

```
mixtwo(n, sigma, binary)
```

Arguments

n	The number of observations.
sigma	The error variance.
binary	A boolean argument: <code>binary = TRUE</code> indicates that binary responses are generated and <code>binary = FALSE</code> indicates that continuous responses are generated.

Details

Sample the predictors x_1, \dots, x_{20} from Bernoulli(0.2) independently, x_{21}, \dots, x_{40} from Bernoulli(0.5) independently, and x_{41}, \dots, x_{84} from a multivariate normal distribution with mean 0, variance 1 and correlation 0.3. If `binary = FALSE`, sample the continuous response y from $\text{Normal}(f_0(x), \sigma^2)$, where

$$f_0(x) = -4 + x_1 + \sin(\pi x_1 * x_{44}) - x_{21} + 0.6x_{41} * x_{42} - \exp[-2(x_{42} + 1)^2] - x_{43}^2 + 0.5x_{44}.$$

If `binary = TRUE`, sample the binary response y from $\text{Bernoulli}(\Phi(f_0(x)))$ where f_0 is defined above and Φ is the cumulative density function of the standard normal distribution.

Value

Return a list with the following components.

X	An n by p data frame representing predictors values, with each row corresponding an observation.
Y	A vector of length n representing response values.
f0	A vector of length n representing the values of $f_0(x)$.
sigma	The error variance which is only returned when <code>binary = FALSE</code> .
prob	A vector of length n representing the values of $\Phi(f_0(x))$, which is only returned when <code>binary = TRUE</code> .

Author(s)

Chuji Luo: <cjluo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.

Examples

```
data = mixtwo(100, 1, FALSE)
```

pbart

Probit BART for binary responses with Normal latents

Description

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a binary response y and a p -dimensional vector of predictors $x = (x_1, \dots, x_p)'$, probit BART models y and x using

$$P(Y = 1|x) = \Phi[f(x)],$$

where Φ is the CDF of the standard normal distribution and f is a sum of Bayesian regression trees function.

The function `pbart()` is inherited from the CRAN R package 'BART' and two modifications are made for the splitting probability and variable importance (see Details).

Usage

```
pbart(
  x.train,
  y.train,
  x.test = matrix(0, 0, 0),
  sparse = FALSE,
  theta = 0,
  omega = 1,
  a = 0.5,
  b = 1,
  augment = FALSE,
  rho = NULL,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
  rm.const = TRUE,
  grp = NULL,
```

```

xnames = NULL,
categorical.idx = NULL,
k = 2,
power = 2,
base = -1,
split.prob = "polynomial",
binaryOffset = NULL,
ntree = 50L,
ndpost = 1000L,
nskip = 1000L,
keepevery = 1L,
nkeeptrain = ndpost,
nkeeptest = ndpost,
nkeepreedraws = ndpost,
printevery = 100L,
transposed = FALSE,
verbose = FALSE
)

```

Arguments

<code>x.train</code>	A matrix or a data frame of predictors values (for training) with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y.train</code>	A vector of continuous response values for training.
<code>x.test</code>	A matrix or a data frame of predictors values for testing, which has the same structure as <code>x.train</code> .
<code>sparse</code>	A Boolean argument indicating whether to replace the discrete uniform distribution for selecting a split variable with a categorical distribution whose event probabilities follow a Dirichlet distribution (see Linero (2018) for details).
<code>theta</code>	Set theta parameter; zero means random.
<code>omega</code>	Set omega parameter; zero means random.
<code>a</code>	A sparse parameter of $Beta(a, b)$ hyper-prior where $0.5 \leq a \leq 1$; a lower value induces more sparsity.
<code>b</code>	A sparse parameter of $Beta(a, b)$ hyper-prior; typically, $b = 1$.
<code>augment</code>	A Boolean argument indicating whether data augmentation is performed in the variable selection procedure of Linero (2018).
<code>rho</code>	A sparse parameter; typically $\rho = p$ where p is the number of predictors.
<code>xinfo</code>	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
<code>numcut</code>	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x.train)</code> is required, where the i -th element gives the number of cut-points for the i -th

	predictor in <code>x.train</code> . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x.train</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x.train - 1)</code> cut-point values are used.
<code>usequants</code>	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
<code>cont</code>	A Boolean argument indicating whether to assume all predictors are continuous.
<code>rm.const</code>	A Boolean argument indicating whether to remove constant predictors.
<code>grp</code>	A vector of group indices for predictors. For example, if 2 appears 3 times in <code>grp</code> , the second predictor of <code>x.train</code> is a categorical predictor with 3 levels. <code>grp</code> is required if <code>transposed=TRUE</code> .
<code>xnames</code>	Column names of <code>x.train</code> . <code>xnames</code> is required if <code>transposed=TRUE</code> .
<code>categorical.idx</code>	A vector of the column indices of categorical predictors in <code>x.train</code> . <code>categorical.idx</code> is required if <code>transposed=TRUE</code> .
<code>k</code>	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y.train</code>) is internally scaled to the range from $-.5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
<code>power</code>	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
<code>base</code>	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
<code>split.prob</code>	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
<code>binaryOffset</code>	The binary offset term in the probit BART model, i.e., $P(Y = 1 x) = \Phi[f(x) + \text{binaryOffset}]$.
<code>ntree</code>	The number of trees in the ensemble.
<code>ndpost</code>	The number of posterior samples returned.
<code>nskip</code>	The number of posterior samples burned in.
<code>keepevery</code>	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
<code>nkeeptrain</code>	The number of posterior samples returned for the train data.
<code>nkeeptest</code>	The number of posterior samples returned for the test data.
<code>nkeeptreedraws</code>	The number of posterior samples returned for the tree draws.
<code>printevery</code>	As the MCMC runs, a message is printed every <code>printevery</code> iterations.
<code>transposed</code>	A Boolean argument indicating whether the matrices <code>x.train</code> and <code>x.test</code> are transposed.
<code>verbose</code>	A Boolean argument indicating whether any messages are printed out.

Details

This function is inherited from `BART::pbart()`. While the original features of `BART::pbart()` are preserved, two modifications are made.

The first modification is to provide two types of split probability for BART. One split probability is proposed in Chipman et al. (2010) and defined as

$$p(d) = \gamma * (1 + d)^{-\beta},$$

where d is the depth of the node, $\gamma \in (0, 1)$ and $\beta \in (0, \infty)$. The other split probability is proposed by Rockova and Saha (2019) and defined as

$$p(d) = \gamma^d,$$

where $\gamma \in (1/n, 1/2)$. BART with the second split probability is proved to achieve the optimal posterior contraction.

The second modification is to provide five types of variable importance measures (`vip`, `within.type.vip`, `pvip`, `varprob.mean` and `mi`) in the return object, for the sake of the existence of mixed-type predictors.

Value

The function `pbart()` returns an object of type `pbart` which essentially is a list consisting of the following components.

<code>yhat.train</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.train)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a training data point. The (i, j) -th element of the matrix is $f^*(x) + \text{binaryOffset}$ for the i -th kept draw of f and the j -th training data point. Burn-in posterior samples are dropped.
<code>prob.train</code>	A matrix with the same structure as <code>yhat.train</code> and the (i, j) -th element of <code>prob.train</code> is $\Phi[f^*(x) + \text{binaryOffset}]$ for the i -th kept draw of f and the j -th training data point.
<code>prob.train.mean</code>	A vector which is <code>colMeans(prob.train)</code> .
<code>yhat.test</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.test)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a test data point. The (i, j) -th element of the matrix is $f^*(x) + \text{binaryOffset}$ for the i -th kept draw of f and the j -th test data point. Burn-in posterior samples are dropped.
<code>prob.test</code>	A matrix with the same structure as <code>yhat.test</code> and the (i, j) -th element of <code>prob.test</code> is $\Phi[f^*(x) + \text{binaryOffset}]$ for the i -th kept draw of f and the j -th test data point.
<code>prob.test.mean</code>	A vector which is <code>colMeans(prob.test)</code> .
<code>varcount</code>	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the number of times that the j -th predictor is used as a split variable in the i -th posterior sample.

varprob	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the split probability of the j -th predictor in the i -th posterior sample. Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
treedraws	A list containing the posterior samples of the ensembles (trees structures, split variables and split values); Can be used for prediction.
proc.time	The process time of running the function <code>wbart()</code> .
vip	A vector of variable inclusion proportions (VIP) proposed in Chipman et al. (2010).
within.type.vip	A vector of within-type VIPs proposed in Luo and Daniels (2021).
pvip	A vector of marginal posterior variable inclusion probabilities (PVIP) proposed in Linero (2018); Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
varprob.mean	A vector of posterior split probabilities (PSP) proposed in Linero (2018); Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
mr.vecs	A list of <code>ncol(x.train)</code> sub-lists with each corresponding to a predictor; Each sub-list contains <code>ndpost</code> vectors with each vector containing the (birth) Metropolis ratios for splits using the predictor as the split variable in that posterior sample.
mr.mean	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the average Metropolis acceptance ratio per splitting rule using the j -th predictor in the i -th posterior sample.
mi	A vector of Metropolis importance (MI) proposed in Luo and Daniels (2021).
rm.const	A vector of indicators for the predictors (after dummification) used in BART; when the indicator is negative, it refers to remove that predictor.
binaryOffset	The binary offset term used in BART.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[mc.pbart](#), [wbart](#) and [pwbart](#).

Examples

```
## simulate data (Scenario B.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, TRUE)
## test pbart() function
res = pbart(data$X, data$Y, ntree=10, nskip=100, ndpost=100)
```

permute.vs

Permutation-based variable selection approach

Description

This function implements the permutation-based variable selection approach for BART (see Algorithm 1 in Luo and Daniels (2021) for details). Three types of variable importance measures are considered: BART variable inclusion proportions (VIP), BART within-type variable inclusion proportions (within-type VIP) and BART Metropolis Importance (MI).

The permutation-based variable selection approach using BART VIP as the variable importance measure is proposed by Bleich et al. (2014). BART within-type VIP and BART MI are proposed by Luo and Daniels (2021), for the sake of the existence of mixed-type predictors and the goal of allowing more relevant predictors into the model.

Usage

```
permute.vs(
  x.train,
  y.train,
  probit = FALSE,
  npermute = 100L,
  nreps = 10L,
  alpha = 0.05,
  true.idx = NULL,
  plot = TRUE,
  n.var.plot = Inf,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
  rm.const = TRUE,
  k = 2,
  power = 2,
  base = 0.95,
  split.prob = "polynomial",
```



```

ntree = 20L,
ndpost = 1000,
nskip = 1000,
keepevery = 1L,
printevery = 100L,
verbose = FALSE
)

```

Arguments

<code>x.train</code>	A matrix or a data frame of predictors values with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y.train</code>	A vector of response (continuous or binary) values.
<code>probit</code>	A Boolean argument indicating whether the response variable is binary or continuous; <code>probit=FALSE</code> (by default) means that the response variable is continuous.
<code>npermute</code>	The number of permutations for estimating the null distributions of the variable importance scores.
<code>nreps</code>	The number of replications for obtaining the averaged (or median) variable importance scores based on the original data set.
<code>alpha</code>	A number between 0 and 1; a predictor is selected if its averaged (or median) variable importance score exceeds the $1 - \alpha$ quantile of the corresponding null distribution.
<code>true.idx</code>	(Optional) A vector of indices of the true relevant predictors; if provided, metrics including precision, recall and F1 score will be returned.
<code>plot</code>	(Optional) A Boolean argument indicating whether plots are returned or not.
<code>n.var.plot</code>	The number of variables to be plotted.
<code>xinfo</code>	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
<code>numcut</code>	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x.train)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in <code>x.train</code> . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x.train</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x.train - 1)</code> cut-point values are used.
<code>usequants</code>	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
<code>cont</code>	A Boolean argument indicating whether to assume all predictors are continuous.
<code>rm.const</code>	A Boolean argument indicating whether to remove constant predictors.

<code>k</code>	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y.train</code>) is internally scaled to the range from $-.5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
<code>power</code>	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
<code>base</code>	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is <code>base^d</code> .
<code>split.prob</code>	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
<code>nree</code>	The number of trees in the ensemble.
<code>ndpost</code>	The number of posterior samples returned.
<code>nskip</code>	The number of posterior samples burned in.
<code>keepevery</code>	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
<code>printevery</code>	As the MCMC runs, a message is printed every <code>printevery</code> iterations.
<code>verbose</code>	A Boolean argument indicating whether any messages are printed out.

Details

The detailed algorithm can be found in Algorithm 1 in Luo and Daniels (2021). The permutation-based variable selection approach using within-type VIP as the variable importance measure is only used when the predictors are of mixed-type; otherwise, it is the same as the one using VIP as the variable importance measure.

If `true.idx` is provided, the precision, recall and F1 scores will be returned for the three (or two if the predictors are of the same type) methods.

If `plot=TRUE`, three (or two if the predictors are of the same type) plots showing which predictors are selected are generated.

Value

The function `permute.vs()` returns three (or two if the predictors are of the same type) plots if `plot=TRUE` and a list with the following components.

<code>vip.imp.cols</code>	The vector of column indices of the predictors selected by the approach using VIP as the variable importance score.
<code>vip.imp.names</code>	The vector of column names of the predictors selected by the approach using VIP as the variable importance score.
<code>avg.vip</code>	The vector (length= <code>ncol(x.train)</code>) of the averaged VIPs based on the original data set; <code>avg.vip=colMeans(avg.vip.mtx)</code> .
<code>avg.vip.mtx</code>	A matrix of VIPs based on the original data set, with each row corresponding to a repetition and each column corresponding to a predictor.
<code>permute.vips</code>	A matrix of VIPs based on the null data sets, with each row corresponding to a permutation (null data set) and each column corresponding to a predictor.

<code>within.type.vip.imp.cols</code>	The vector of column indices of the predictors selected by the approach using within-type VIP as the variable importance score.
<code>within.type.vip.imp.names</code>	The vector of column names of the predictors selected by the approach using within-type VIP as the variable importance score.
<code>avg.within.type.vip</code>	The vector (<code>length=ncol(x.train)</code>) of the averaged within-type VIPs based on the original data set; <code>avg.within.type.vip=colMeans(avg.within.type.vip.mtx)</code> .
<code>avg.within.type.vip.mtx</code>	A matrix of within-type VIPs based on the original data set, with each row corresponding to a repetition and each column corresponding to a predictor.
<code>permutate.within.type.vips</code>	A matrix of within VIPs based on the null data sets, with each row corresponding to a permutation (null data set) and each column corresponding to a predictor.
<code>varcounts</code>	A list of <code>nreps+npermutate</code> elements; Each element is a <code>ndpost</code> by <code>ncol(x.train)</code> matrix with each row corresponding to a draw of the ensemble and each column corresponding to a predictor; The (i, j) -th element of the matrix is the number of times that the j -th predictor is used as a split variable in the i -th posterior sample; The first <code>nreps</code> elements correspond to <code>nreps</code> repetitions and the latter <code>npermutate</code> elements correspond to <code>npermutate</code> permutations.
<code>mi.imp.cols</code>	The vector of column indices of the predictors selected by the approach using MI as the variable importance score.
<code>mi.imp.names</code>	The vector of column names of the predictors selected by the approach using MI as the variable importance score.
<code>median.mi</code>	The vector (<code>length=ncol(x.train)</code>) of the median MIs based on the original data set; <code>median.mi=colMeans(median.mi.mtx)</code> .
<code>median.mi.mtx</code>	A matrix of MIs based on the original data set, with each row corresponding to a repetition and each column corresponding to a predictor.
<code>permutate.mis</code>	A matrix of MIs based on the null data sets, with each row corresponding to a permutation (null data set) and each column corresponding to a predictor.
<code>true.idx</code>	A vector of indices of the true relevant predictors; only returned if <code>true.idx</code> is provided as inputs.
<code>vip.precision</code>	The precision score for the approach using VIP as the variable importance score; only returned if <code>true.idx</code> is provided.
<code>vip.recall</code>	The recall score for the approach using VIP as the variable importance score; only returned if <code>true.idx</code> is provided.
<code>vip.f1</code>	The F1 score for the approach using VIP as the variable importance score; only returned if <code>true.idx</code> is provided.
<code>wt.vip.precision</code>	The precision score for the approach using within-VIP as the variable importance score; only returned when the predictors are of the same type and <code>true.idx</code> is provided.

wt.vip.recall	The recall score for the approach using within-VIP as the variable importance score; only returned when the predictors are of the same type and true.idx is provided.
wt.vip.f1	The F1 score for the approach using within-VIP as the variable importance score; only returned when the predictors are of the same type and true.idx is provided.
mi.precision	The precision score for the approach using MI as the variable importance score; only returned if true.idx is provided.
mi.recall	The recall score for the approach using MI as the variable importance score; only returned if true.idx is provided.
mi.f1	The F1 score for the approach using MI as the variable importance score; only returned if true.idx is provided.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Bleich, Justin et al. (2014). "Variable selection for BART: an application to gene regulation." *Ann. Appl. Stat.* 8.3, pp 1750–1781.
- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* 4 266–298.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.

See Also

[mc.permute.vs](#), [medianInclusion.vs](#), [mc.backward.vs](#) and [abc.vs](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## test permute.vs() function
res = permute.vs(data$X, data$Y, probit=FALSE, npermute=100, nreps=10, alpha=0.05,
true.idx=c(1,2,6:8), plot=FALSE, ntree=10, ndpost=100, nskip=100)
```

predict.pbart

*Predict new observations with a fitted BART model***Description**

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a binary response y , probit BART models y and x using

$$P(Y = 1|x) = \Phi[f(x)],$$

where Φ is the CDF of the standard normal distribution and f is a sum of Bayesian regression trees function.

This function uses S3 method for the class `pbart` and is inherited from the CRAN R package 'BART'.

Usage

```
## S3 method for class 'pbart'
predict(object, newdata, mc.cores = 1, openmp = (mc.cores.openmp() > 0), ...)
```

Arguments

<code>object</code>	An object of class <code>pbart</code> , returned from the function <code>pbart()</code> .
<code>newdata</code>	A matrix of predictors with rows corresponding to new observations.
<code>mc.cores</code>	The number of threads to utilize.
<code>openmp</code>	A Boolean argument dictating whether OpenMP is utilized for parallel processing. This depends on whether OpenMP is available on your system which, by default, is verified with the function <code>mc.cores.openmp()</code> .
<code>...</code>	Other arguments passed on to the function <code>pwbart()</code> .

Value

Returns a matrix of prediction for `newdata`, whose rows correspond to draws and columns correspond to observations.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.

Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.

Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.

Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.

Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[pwbart](#) and [pbart](#).

Examples

```
## simulate data (Scenario B.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, TRUE)
## run pbart() function
res = pbart(data$X, data$Y, ntree=10, nskip=100, ndpost=100)
## test predict.pbart() function
newdata = mixone(5, 10, 1, TRUE)$X
pred = predict(res, newdata)
```

predict.wbart

Predict new observations with a fitted BART model

Description

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a continuous response y and a p -dimensional vector of predictors $x = (x_1, \dots, x_p)'$, BART models y and x using

$$y = f(x) + \epsilon,$$

where f is a sum of Bayesian regression trees function and $\epsilon \sim N(0, \sigma^2)$.

This function uses S3 method for the class `wbart` and is inherited from the CRAN R package 'BART'.

Usage

```
## S3 method for class 'wbart'
predict(object, newdata, mc.cores = 1, openmp = (mc.cores.openmp() > 0), ...)
```

Arguments

object	An object of class <code>wbart</code> , returned from the function <code>wbart()</code> .
newdata	A matrix of predictors with rows corresponding to new observations.
mc.cores	The number of threads to utilize.
openmp	A Boolean argument dictating whether OpenMP is utilized for parallel processing. This depends on whether OpenMP is available on your system which, by default, is verified with the function <code>mc.cores.openmp()</code> .
...	Other arguments passed on to the function <code>pwbart()</code> .

Value

Returns a matrix of prediction for `newdata`, whose rows correspond to draws and columns correspond to observations.

Author(s)

Chuji Luo: <cjl@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[pwbart](#) and [wbart](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## run wbart() function
res = wbart(data$X, data$Y, ntree=10, nskip=100, ndpost=100)
## test predict.wbart() function
newdata = mixone(5, 10, 1, FALSE)$X
pred = predict(res, newdata)
```

pwbart

*Predicting new observations with a previously fitted BART model***Description**

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a continuous response y and a p -dimensional vector of predictors $x = (x_1, \dots, x_p)'$, BART models y and x using

$$y = f(x) + \epsilon,$$

where f is a sum of Bayesian regression trees function and $\epsilon \sim N(0, \sigma^2)$.

For a binary response y , probit BART models y and x using

$$P(Y = 1|x) = \Phi[f(x)],$$

where Φ is the CDF of the standard normal distribution and f is a sum of Bayesian regression trees function.

The function `pwbart()` is inherited from the CRAN R package 'BART'.

Usage

```
pwbart(
  x.test,
  treedraws,
  rm.const,
  mu = 0,
  mc.cores = 1L,
  transposed = FALSE,
  dodraws = TRUE,
  verbose = FALSE
)
```

Arguments

<code>x.test</code>	A matrix or a data frame of predictors values for prediction with each row corresponding to an observation and each column corresponding to a predictor.
<code>treedraws</code>	A list which is the <code>\$treedraws</code> returned from the function <code>wbart()</code> or <code>pbart()</code> .
<code>rm.const</code>	A vector which is the <code>\$rm.const</code> returned from the function <code>wbart()</code> or <code>pbart()</code> .
<code>mu</code>	Mean to add on to y prediction.
<code>mc.cores</code>	The number of threads to utilize.
<code>transposed</code>	A Boolean argument indicating whether the matrix <code>x.test</code> is transposed. When running <code>pwbart()</code> or <code>mc.pwbart()</code> in parallel, it is more memory-efficient to transpose <code>x.test</code> prior to calling the internal versions of these functions.
<code>dodraws</code>	A Boolean argument indicating whether to return the draws themselves (the default), or whether to return the mean of the draws as specified by <code>dodraws=FALSE</code> .
<code>verbose</code>	A Boolean argument indicating whether any messages are printed out.

Value

Returns the predictions for `x.test`. If `dodraws=TRUE`, return a matrix of prediction with each row corresponding to a draw and each column corresponding to a new observation; if `dodraws=FALSE`, return a vector of predictions which are the mean of the draws.

Author(s)

Chuji Luo: <cjluo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[wbart](#), [pbart](#) and [mc.pwbart](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## run wbart() function
res = wbart(data$X, data$Y, ntree=10, nskip=100, ndpost=100)
## test pwbart() function
x.test = mixone(5, 10, 1, FALSE)$X
pred = pwbart(x.test, res$treedraws, res$rm.const, mu=mean(data$Y))
```

Description

BART is a Bayesian approach to nonparametric function estimation and inference using a sum of trees.

For a continuous response y and a p -dimensional vector of predictors $x = (x_1, \dots, x_p)'$, BART models y and x using

$$y = f(x) + \epsilon,$$

where f is a sum of Bayesian regression trees function and $\epsilon \sim N(0, \sigma^2)$.

The function `wbart()` is inherited from the CRAN R package 'BART' and two modifications are made for the splitting probability and variable importance (see Details).

Usage

```
wbart(
  x.train,
  y.train,
  x.test = matrix(0, 0, 0),
  sparse = FALSE,
  theta = 0,
  omega = 1,
  a = 0.5,
  b = 1,
  augment = FALSE,
  rho = NULL,
  xinfo = matrix(0, 0, 0),
  numcut = 100L,
  usequants = FALSE,
  cont = FALSE,
  rm.const = TRUE,
  grp = NULL,
  xnames = NULL,
  categorical.idx = NULL,
  power = 2,
  base = -1,
  split.prob = "polynomial",
  k = 2,
  sigmaf = NA,
  sigest = NA,
  sigdf = 3,
  sigquant = 0.9,
  lambda = NA,
  fmean = mean(y.train),
  w = rep(1, length(y.train)),
  ntree = 200L,
  ndpost = 1000L,
  nskip = 1000L,
  keepevery = 1L,
  nkeeptrain = ndpost,
  nkeptest = ndpost,
```

```

nkeepstestmean = ndpost,
nkeepreedraws = ndpost,
printevery = 100L,
transposed = FALSE,
verbose = FALSE
)

```

Arguments

<code>x.train</code>	A matrix or a data frame of predictors values (for training) with each row corresponding to an observation and each column corresponding to a predictor. If a predictor is a factor with q levels in a data frame, it is replaced with q dummy variables.
<code>y.train</code>	A vector of continuous response values for training.
<code>x.test</code>	A matrix or a data frame of predictors values for testing, which has the same structure as <code>x.train</code> .
<code>sparse</code>	A Boolean argument indicating whether to replace the discrete uniform distribution for selecting a split variable with a categorical distribution whose event probabilities follow a Dirichlet distribution (see Linero (2018) for details).
<code>theta</code>	Set theta parameter; zero means random.
<code>omega</code>	Set omega parameter; zero means random.
<code>a</code>	A sparse parameter of $Beta(a, b)$ hyper-prior where $0.5 \leq a \leq 1$; a lower value induces more sparsity.
<code>b</code>	A sparse parameter of $Beta(a, b)$ hyper-prior; typically, $b = 1$.
<code>augment</code>	A Boolean argument indicating whether data augmentation is performed in the variable selection procedure of Linero (2018).
<code>rho</code>	A sparse parameter; typically $\rho = p$ where p is the number of predictors.
<code>xinfo</code>	A matrix of cut-points with each row corresponding to a predictor and each column corresponding to a cut-point. <code>xinfo=matrix(0.0,0,0)</code> indicates the cut-points are specified by BART.
<code>numcut</code>	The number of possible cut-points; If a single number is given, this is used for all predictors; Otherwise a vector with length equal to <code>ncol(x.train)</code> is required, where the i -th element gives the number of cut-points for the i -th predictor in <code>x.train</code> . If <code>usequants=FALSE</code> , <code>numcut</code> equally spaced cut-points are used to cover the range of values in the corresponding column of <code>x.train</code> . If <code>usequants=TRUE</code> , then <code>min(numcut, the number of unique values in the corresponding column of x.train - 1)</code> cut-point values are used.
<code>usequants</code>	A Boolean argument indicating how the cut-points in <code>xinfo</code> are generated; If <code>usequants=TRUE</code> , uniform quantiles are used for the cut-points; Otherwise, the cut-points are generated uniformly.
<code>cont</code>	A Boolean argument indicating whether to assume all predictors are continuous.
<code>rm.const</code>	A Boolean argument indicating whether to remove constant predictors.
<code>grp</code>	A vector of group indices for predictors. For example, if 2 appears 3 times in <code>grp</code> , the second predictor of <code>x.train</code> is a categorical predictor with 3 levels. <code>grp</code> is required if <code>transposed=TRUE</code> .

xnames	Column names of <code>x.train</code> . <code>xnames</code> is required if <code>transposed=TRUE</code> .
categorical.idx	A vector of the column indices of categorical predictors in <code>x.train</code> . <code>categorical.idx</code> is required if <code>transposed=TRUE</code> .
power	The power parameter of the polynomial splitting probability for the tree prior. Only used if <code>split.prob="polynomial"</code> .
base	The base parameter of the polynomial splitting probability for the tree prior if <code>split.prob="polynomial"</code> ; if <code>split.prob="exponential"</code> , the probability of splitting a node at depth d is base^d .
split.prob	A string indicating what kind of splitting probability is used for the tree prior. If <code>split.prob="polynomial"</code> , the splitting probability in Chipman et al. (2010) is used; If <code>split.prob="exponential"</code> , the splitting probability in Rockova and Saha (2019) is used.
k	The number of prior standard deviations that $E(Y x) = f(x)$ is away from $+/- .5$. The response (<code>y.train</code>) is internally scaled to the range from $- .5$ to $.5$. The bigger <code>k</code> is, the more conservative the fitting will be.
sigmaf	The standard deviation of <code>f</code> .
sigest	A rough estimate of the error standard deviation, the square of which follows an inverse chi-squared prior. If <code>sigest=NA</code> , the rough estimate will be the usual least square estimator; Otherwise, the supplied value will be used.
sigdf	The degrees of freedom for the error variance prior.
sigquant	The quantile of the error variance prior, where <code>sigest</code> is placed. The closer the quantile is to 1, the more aggressive the fit will be.
lambda	The scale parameter of the error variance prior.
fmean	BART operates on <code>y.train</code> centered by <code>fmean</code> .
w	A vector of weights which multiply the standard deviation.
ntree	The number of trees in the ensemble.
ndpost	The number of posterior samples returned.
nskip	The number of posterior samples burned in.
keepevery	Every <code>keepevery</code> posterior sample is kept to be returned to the user.
nkeeptrain	The number of posterior samples returned for the train data.
nkeptest	The number of posterior samples returned for the test data.
nkeptestmean	The number of posterior samples returned for the test mean.
nkeptreedraws	The number of posterior samples returned for the tree draws.
printevery	As the MCMC runs, a message is printed every <code>printevery</code> iterations.
transposed	A Boolean argument indicating whether the matrices <code>x.train</code> and <code>x.test</code> are transposed.
verbose	A Boolean argument indicating whether any messages are printed out.

Details

This function is inherited from `BART::wbart()`. While the original features of `BART::wbart()` are preserved, two modifications are made.

The first modification is to provide two types of split probability for BART. One split probability is proposed in Chipman et al. (2010) and defined as

$$p(d) = \gamma * (1 + d)^{-\beta},$$

where d is the depth of the node, $\gamma \in (0, 1)$ and $\beta \in (0, \infty)$. The other split probability is proposed by Rockova and Saha (2019) and defined as

$$p(d) = \gamma^d,$$

where $\gamma \in (1/n, 1/2)$. BART with the second split probability is proved to achieve the optimal posterior contraction.

The second modification is to provide five types of variable importance measures (`vip`, `within.type.vip`, `pvip`, `varprob.mean` and `mi`) in the return object, for the sake of the existence of mixed-type predictors.

Value

The function `wbart()` returns an object of type `wbart` which essentially is a list consisting of the following components.

<code>sigma</code>	A vector with <code>nskip+ndpost*keepevery</code> posterior samples of σ .
<code>yhat.train.mean</code>	<code>colMeans(yhat.train)</code> .
<code>yhat.train</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.train)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a training data point. The (i, j) -th element of the matrix is $f^*(x)$ for the i -th kept draw of f and the j -th training data point. Burn-in posterior samples are dropped.
<code>yhat.test.mean</code>	<code>colMeans(yhat.test)</code> .
<code>yhat.test</code>	A matrix with <code>ndpost</code> rows and <code>nrow(x.test)</code> columns with each row corresponding to a draw f^* from the posterior of f and each column corresponding to a test data point. The (i, j) -th element of the matrix is $f^*(x)$ for the i -th kept draw of f and the j -th test data point. Burn-in posterior samples are dropped.
<code>varcount</code>	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the number of times that the j -th predictor is used as a split variable in the i -th posterior sample.
<code>varprob</code>	A matrix with <code>ndpost</code> rows and <code>ncol(x.train)</code> columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the split probability of the j -th predictor in the i -th posterior sample. Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
<code>treedraws</code>	A list containing the posterior samples of the ensembles (trees structures, split variables and split values); Can be used for prediction.

<code>proc.time</code>	The process time of running the function <code>wbart()</code> .
<code>mu</code>	BART operates on <code>y.train</code> centered by <code>fmean</code> .
<code>mr.vecs</code>	A list of $n_{col}(x.train)$ sub-lists with each corresponding to a predictor; Each sub-list contains <code>ndpost</code> vectors with each vector containing the (birth) Metropolis ratios for splits using the predictor as the split variable in that posterior sample.
<code>vip</code>	A vector of variable inclusion proportions (VIP) proposed in Chipman et al. (2010).
<code>within.type.vip</code>	A vector of within-type VIPs proposed in Luo and Daniels (2021).
<code>pvip</code>	A vector of marginal posterior variable inclusion probabilities (PVIP) proposed in Linero (2018); Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
<code>varprob.mean</code>	A vector of posterior split probabilities (PSP) proposed in Linero (2018); Only useful when DART is fit, i.e., <code>sparse=TRUE</code> .
<code>mr.mean</code>	A matrix with <code>ndpost</code> rows and $n_{col}(x.train)$ columns with each row corresponding to a draw of the ensemble and each column corresponding to a predictor. The (i, j) -th element is the average Metropolis acceptance ratio per splitting rule using the j -th predictor in the i -th posterior sample.
<code>mi</code>	A vector of Metropolis importance (MI) proposed in Luo and Daniels (2021).
<code>rm.const</code>	A vector of indicators for the predictors (after dummification) used in BART; when the indicator is negative, it refers to remove that predictor.

Author(s)

Chuji Luo: <cjl1uo@ufl.edu> and Michael J. Daniels: <daniels@ufl.edu>.

References

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2010). "BART: Bayesian additive regression trees." *Ann. Appl. Stat.* **4** 266–298.
- Linero, A. R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection." *J. Amer. Statist. Assoc.* **113** 626–636.
- Luo, C. and Daniels, M. J. (2021) "Variable Selection Using Bayesian Additive Regression Trees." *arXiv preprint arXiv:2112.13998*.
- Rockova V, Saha E (2019). "On theory for BART." *In The 22nd International Conference on Artificial Intelligence and Statistics* (pp. 2839–2848). PMLR.
- Sparapani, R., Spanbauer, C. and McCulloch, R. (2021). "Nonparametric machine learning and efficient computation with bayesian additive regression trees: the BART R package." *J. Stat. Softw.* **97** 1–66.

See Also

[mc.wbart](#), [pbart](#) and [pwbart](#).

Examples

```
## simulate data (Scenario C.M.1. in Luo and Daniels (2021))
set.seed(123)
data = mixone(100, 10, 1, FALSE)
## test wbart() function
res = wbart(data$X, data$Y, ntree=10, nskip=100, ndpost=100)
```

Index

* **BART**

BartMixVs-package, 2

* **Non-Parametric Regression**

BartMixVs-package, 2

* **Variable Selection**

BartMixVs-package, 2

abc.vs, 3, 15, 19, 30, 40, 52

BartMixVs (BartMixVs-package), 2

BartMixVs-package, 2

bartModelMatrix, 7

checkerboard, 9

friedman, 10

mc.abc.vs, 11

mc.backward.vs, 7, 16, 30, 40, 52

mc.cores.openmp, 19

mc.pbart, 20, 48

mc.permute.vs, 25, 52

mc.pwbart, 30, 31, 57

mc.wbart, 32, 62

medianInclusion.vs, 7, 19, 30, 37, 52

mixone, 41

mixtwo, 42

pbart, 9, 25, 31, 43, 54, 57, 62

permute.vs, 7, 19, 30, 40, 48

predict.pbart, 53

predict.wbart, 54

pwbart, 20, 48, 54, 55, 56, 62

wbart, 9, 31, 37, 48, 55, 57, 57