

Package ‘McMiso’

April 3, 2026

Type Package

Title Multicore Multivariable Isotonic Regression

Version 0.2.0

Description Provides functions for isotonic regression and classification when there are multiple independent variables. The functions solve the optimization problem using a projective Bayes approach with recursive sequential update algorithms, and are useful for situations with a relatively large number of covariates. Supports binary outcomes via a Beta-Binomial conjugate model ('miso', 'PBclassifier') and continuous outcomes via a Normal-Inverse-Chi-Squared conjugate model ('misoN'). Parallel computing wrappers ('mcmiso', 'mcPBclassifier', 'mcmisoN') are provided that run the down-up and up-down algorithms simultaneously and return whichever finishes first. The estimation method follows the projective Bayes solution described in Cheung and Diaz (2023) [doi:10.1093/jrsssb/qkad014](https://doi.org/10.1093/jrsssb/qkad014).

Depends R (>= 4.0.0)

Imports stats, utils

Suggests future (>= 1.33.0)

License GPL-3

Encoding UTF-8

RoxygenNote 7.3.2

NeedsCompilation no

Author Cheung Ken [aut, cre]

Maintainer Cheung Ken <yc632@cumc.columbia.edu>

Repository CRAN

Date/Publication 2026-04-03 17:40:02 UTC

Contents

boundary	2
mcmiso	3

mcmisoN	4
mcPBclassifier	6
miso	8
misoN	9
PBclassifier	11
predict.pbc	12

Index	14
--------------	-----------

boundary	<i>Decision Boundary of a Projective Bayes Classifier</i>
----------	---

Description

Extracts the decision boundary from a fitted projective Bayes classifier. The decision boundary consists of the minimal set of feature combinations that are classified as positive (1), i.e., combinations classified as 1 for which no combination lower in the partial ordering is also classified as 1. Any feature combination at or above a boundary combination in the partial ordering is guaranteed to be classified as 1, making the boundary the most compact representation of the classification rule.

Usage

```
boundary(object)
```

Arguments

object a fitted object of class "pbc", produced by [PBclassifier](#) or [mcPBclassifier](#).

Value

A list of class "boundary" containing the following components:

boundary a numeric matrix of feature combinations forming the decision boundary, where each row is one minimal positive combination. NULL if no combinations are classified as 1.

yhat a binary numeric vector giving the classification of each boundary combination (all values will be 1)

pt a numeric vector of posterior probabilities at each boundary combination

n_boundary an integer giving the number of boundary combinations

n_positive an integer giving the total number of combinations classified as 1

n_total an integer giving the total number of unique feature combinations

References

Cheung YK, Diaz KM. Monotone response surface of multi-factor condition: estimation and Bayes classifiers. *J R Stat Soc Series B Stat Methodol.* 2023 Apr;85(2):497-522. doi: 10.1093/jrsssb/qkad014. Epub 2023 Mar 22. PMID: 38464683; PMCID: PMC10919322.

Cheung YK, Kuhn L. Evaluating multiplex diagnostic test using partially ordered Bayes classifier. *Ann Appl Stat.* In press.

Examples

```
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A), size=500, replace=TRUE),]
y <- as.numeric(rowSums(X) >= 3)
fit <- PBclassifier(X, y)
db <- boundary(fit)
print(db)
```

mcmiso

Multivariable Isotonic Regression for Binary Data using Inverse Projective Bayes with Parallel Computing

Description

A parallel computing wrapper for `miso` that uses `mcComb` to run the down-up ("DU") and up-down ("UD") algorithms simultaneously at each threshold grid point, returning the result of whichever finishes first. While the two algorithms may produce different classification rules at each threshold, they are guaranteed to achieve the same maximum log posterior gain (Cheung and Kuhn, in press). This approach reduces elapsed time without sacrificing accuracy and is most effective for large datasets where the number of unique feature combinations K is large (typically $K > 500$).

Usage

```
mcmiso(X, y, incr = 0.01)
```

Arguments

<code>X</code>	a numeric matrix of observed feature combinations, one row per observation, where repeated rows are expected. Each column represents a feature (e.g., a dose component or experimental factor) and each row represents the feature combination observed for one unit.
<code>y</code>	a binary numeric vector of length <code>nrow(X)</code> indicating the observed outcome for each observation (1 = event, 0 = no event).
<code>incr</code>	a numeric value in (0,1) specifying the increment between threshold grid points used to invert the classifier. Smaller values yield finer resolution at the cost of increased computation time. Defaults to 0.01.

Details

Before calling this function, the user must set up a parallel plan using `future::plan`. The recommended plan is `future::multisession` which works across all platforms including Windows. The `future` package must be installed separately (`install.packages("future")`). After the analysis is complete, it is good practice to restore the default sequential plan using `future::plan(future::sequential)`.

Note that parallel overhead may outweigh the benefit for small datasets. When run from RStudio, `future::multisession` is used automatically but incurs higher overhead due to session launching costs compared to `future::multicore` available in a terminal.

Value

A list containing the same components as `miso`, based on whichever of the "DU" or "UD" algorithm finishes first at each threshold grid point:

alldoses a numeric matrix of unique feature combinations observed in the training data

M a numeric vector of observation counts at each feature combination

S a numeric vector of event counts at each feature combination

thetahat a numeric vector of estimated response probabilities at each unique feature combination, monotone nondecreasing with respect to the partial ordering of the features

nt an integer giving the number of threshold grid points used, equal to $\text{round}(1/\text{incr}) + 1$

logH a numeric vector of length `nt` giving the log posterior gain of the optimal classification at each threshold grid point

References

Cheung YK, Diaz KM. Monotone response surface of multi-factor condition: estimation and Bayes classifiers. *J R Stat Soc Series B Stat Methodol.* 2023 Apr;85(2):497-522. doi: 10.1093/jrsssb/qkad014. Epub 2023 Mar 22. PMID: 38464683; PMCID: PMC10919322.

Cheung YK, Kuhn L. Evaluating multiplex diagnostic test using partially ordered Bayes classifier. *Ann Appl Stat.* In press.

Examples

```
## Not run:
# install.packages("future") # if not already installed
future::plan(future::multisession) # set up parallel plan first
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A), size=500, replace=TRUE),]
y <- as.numeric(rowSums(X) >= 3)
fit <- mcmiso(X, y)
future::plan(future::sequential) # restore default plan when done

## End(Not run)
```

mcmisoN

Multivariable Isotonic Regression for Continuous Data using Inverse Projective Bayes with Parallel Computing

Description

A parallel computing wrapper for `misoN` that uses `mcComb` to run the down-up ("DU") and up-down ("UD") algorithms simultaneously at each threshold grid point, returning the result of whichever finishes first. While the two algorithms may produce different classification rules at each threshold, they are guaranteed to achieve the same maximum log posterior gain (Cheung and Kuhn, in press). This approach reduces elapsed time without sacrificing accuracy and is most effective for large datasets where the number of unique feature combinations K is large (typically $K > 500$).

Usage

```
mcmisoN(X, y, nt = 101, mu0 = 0, sig0 = 100, kap0 = 0.01, nu0 = 0.01)
```

Arguments

<code>X</code>	a numeric matrix of observed feature combinations, one row per observation, where repeated rows are expected. Each column represents a feature (e.g., a dose component or experimental factor) and each row represents the feature combination observed for one unit.
<code>y</code>	a numeric vector of length <code>nrow(X)</code> containing the continuous outcome for each observation.
<code>nt</code>	a positive integer specifying the number of threshold grid points used to invert the classifier. The grid range is determined automatically from the observed data. Larger values yield finer resolution at the cost of increased computation time. Defaults to 101, which is approximately equivalent to the default <code>incr = 0.01</code> used in <code>mcmiso</code> .
<code>mu0</code>	a numeric value specifying the prior mean of the response. Defaults to 0.
<code>sig0</code>	a positive numeric value specifying the prior scale parameter, interpreted as the prior standard deviation of the response. Defaults to 100, yielding a diffuse prior on the variance.
<code>kap0</code>	a positive numeric value specifying the prior pseudo sample size for the mean. Smaller values yield a more diffuse prior on the mean and reduce prior shrinkage of the posterior mean toward <code>mu0</code> . Defaults to 0.01, approximating a Jeffreys non-informative prior.
<code>nu0</code>	a positive numeric value specifying the prior degrees of freedom for the variance. Smaller values yield a more diffuse prior on the variance. Defaults to 0.01, approximating a Jeffreys non-informative prior.

Details

Before calling this function, the user must set up a parallel plan using `future::plan`. The recommended plan is `future::multisession` which works across all platforms including Windows. The `future` package must be installed separately (`install.packages("future")`). After the analysis is complete, it is good practice to restore the default sequential plan using `future::plan(future::sequential)`.

Note that parallel overhead may outweigh the benefit for small datasets. When run from RStudio, `future::multisession` is used automatically but incurs higher overhead due to session launching costs compared to `future::multicore` available in a terminal.

Value

A list containing the same components as `misoN`, based on whichever of the "DU" or "UD" algorithm finishes first at each threshold grid point:

alldoses a numeric matrix of unique feature combinations observed in the training data

M a numeric vector of observation counts at each feature combination

MY a numeric vector of sample means of the outcome at each feature combination (NA if `M = 0`)

- VY** a numeric vector of sample variances of the outcome at each feature combination (NA if $M \leq 1$)
- thetahat** a numeric vector of estimated mean responses at each unique feature combination, monotone nondecreasing with respect to the partial ordering of the features
- nt** an integer giving the number of threshold grid points used
- logH** a numeric vector of length nt giving the log posterior gain of the optimal classification at each threshold grid point

References

Cheung YK, Diaz KM. Monotone response surface of multi-factor condition: estimation and Bayes classifiers. *J R Stat Soc Series B Stat Methodol.* 2023 Apr;85(2):497-522. doi: 10.1093/jrsssb/qkad014. Epub 2023 Mar 22. PMID: 38464683; PMCID: PMC10919322.

Cheung YK, Kuhn L. Evaluating multiplex diagnostic test using partially ordered Bayes classifier. *Ann Appl Stat.* In press.

Examples

```
## Not run:
# install.packages("future") # if not already installed
future::plan(future::multisession) # set up parallel plan first
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A), size=500, replace=TRUE),]
y <- rowSums(X) + rnorm(500)
fit <- mcmisoN(X, y)
future::plan(future::sequential) # restore default plan when done

## End(Not run)
```

mcPBclassifier

Multivariable Isotonic Classification using Projective Bayes with Parallel Computing

Description

A parallel computing wrapper for [PBclassifier](#) that runs the down-up ("DU") and up-down ("UD") algorithms simultaneously in parallel and returns the result of whichever finishes first. Since both algorithms are guaranteed to achieve the same maximum log posterior gain, i.e. the same $\log H$ (Cheung and Kuhn in press), this approach reduces elapsed time without sacrificing optimality.

Usage

```
mcPBclassifier(X, y, a0 = 0.5, b0 = 0.5, t0 = 0.5)
```

Arguments

<code>X</code>	a numeric matrix of observed feature combinations, one row per observation, where repeated rows are expected. Each column represents a feature (e.g., a dose component or experimental factor) and each row represents the feature combination observed for one unit.
<code>y</code>	a binary numeric vector of length <code>nrow(X)</code> indicating the observed outcome for each observation (1 = event, 0 = no event).
<code>a0</code>	a positive numeric value specifying the <code>shape1</code> hyperparameter of the Beta prior in the Beta-Binomial conjugate model. Defaults to 0.5 (Jeffreys prior).
<code>b0</code>	a positive numeric value specifying the <code>shape2</code> hyperparameter of the Beta prior in the Beta-Binomial conjugate model. Defaults to 0.5 (Jeffreys prior).
<code>t0</code>	a numeric value in (0,1) specifying the threshold on the response rate used to classify each feature combination as event or no-event. Defaults to 0.5.

Details

Before calling this function, the user must set up a parallel plan using `future::plan`. The recommended plan is `future::multisession` which works across all platforms including Windows. The `future` package must be installed separately (`install.packages("future")`). After the analysis is complete, it is good practice to restore the default sequential plan using `future::plan(future::sequential)`.

Note that parallel overhead may outweigh the benefit for small datasets. This function is most effective when the number of unique feature combinations K is large (typically $K > 500$) and when run from a terminal rather than RStudio, where `future::multicore` (forking) is available. In RStudio, `future::multisession` is used automatically but incurs higher overhead due to session launching costs.

Value

A list of class "pbc" containing the same components as `PBclassifier`, based on whichever of the "DU" or "UD" algorithm finishes first.

References

Cheung YK, Diaz KM. Monotone response surface of multi-factor condition: estimation and Bayes classifiers. *J R Stat Soc Series B Stat Methodol.* 2023 Apr;85(2):497-522. doi: 10.1093/jrsssb/qkad014. Epub 2023 Mar 22. PMID: 38464683; PMCID: PMC10919322.

Cheung YK, Kuhn L. Evaluating multiplex diagnostic test using partially ordered Bayes classifier. *Ann Appl Stat.* In press.

Examples

```
## Not run:
# install.packages("future") # if not already installed
future::plan(future::multisession) # set up parallel plan first
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A), size=500, replace=TRUE),]
y <- as.numeric(rowSums(X) >= 3)
```

```
fit <- mcPBclassifier(X, y)
future::plan(future::sequential) # restore default plan when done

## End(Not run)
```

miso *Multivariable Isotonic Regression for Binary Data using Inverse Projective Bayes*

Description

Estimates the underlying response probability for multivariate features using an inverse projective Bayes approach. The estimator is obtained by inverting the projective Bayes classifier across a grid of thresholds, yielding a monotone nonparametric estimate of the response probability that is nondecreasing in each feature. A Beta-Binomial conjugate model is used to compute posterior probabilities at each threshold.

Usage

```
miso(X, y, incr = 0.01)
```

Arguments

X a numeric matrix of observed feature combinations, one row per observation, where repeated rows are expected. Each column represents a feature (e.g., a dose component or experimental factor) and each row represents the feature combination observed for one unit.

y a binary numeric vector of length `nrow(X)` indicating the observed outcome for each observation (1 = event, 0 = no event).

incr a numeric value in (0,1) specifying the increment between threshold grid points used to invert the classifier. Smaller values yield finer resolution at the cost of increased computation time. Defaults to 0.01.

Value

A list containing the following components:

alldoses a numeric matrix of unique feature combinations observed in the training data

M a numeric vector of observation counts at each feature combination

S a numeric vector of event counts at each feature combination

thetahat a numeric vector of estimated response probabilities at each unique feature combination, monotone nondecreasing with respect to the partial ordering of the features

nt an integer giving the number of threshold grid points used, equal to $\text{round}(1/\text{incr}) + 1$

logH a numeric vector of length `nt` giving the log posterior gain of the optimal classification at each threshold grid point

References

Cheung YK, Diaz KM. Monotone response surface of multi-factor condition: estimation and Bayes classifiers. *J R Stat Soc Series B Stat Methodol.* 2023 Apr;85(2):497-522. doi: 10.1093/jrsssb/qkad014. Epub 2023 Mar 22. PMID: 38464683; PMCID: PMC10919322.

Examples

```
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A), size=500, replace=TRUE),]
y <- as.numeric(rowSums(X) >= 3)
miso(X, y)
```

misoN	<i>Multivariable Isotonic Regression for Continuous Data using Inverse Projective Bayes</i>
-------	---

Description

Estimates the underlying mean response for multivariate features using an inverse projective Bayes approach. The estimator is obtained by inverting the projective Bayes classifier across a grid of thresholds, yielding a monotone nonparametric estimate of the mean response that is nondecreasing in each feature. A Normal-Inverse-Chi-Squared conjugate model is used to compute posterior probabilities at each threshold. The threshold grid is determined automatically from the data range and `nt` controls the resolution of the grid.

Usage

```
misoN(X, y, nt = 101, mu0 = 0, sig0 = 100, kap0 = 0.01, nu0 = 0.01)
```

Arguments

<code>X</code>	a numeric matrix of observed feature combinations, one row per observation, where repeated rows are expected. Each column represents a feature (e.g., a dose component or experimental factor) and each row represents the feature combination observed for one unit.
<code>y</code>	a numeric vector of length <code>nrow(X)</code> containing the continuous outcome for each observation.
<code>nt</code>	a positive integer specifying the number of threshold grid points used to invert the classifier. The grid range is determined automatically from the observed data. Larger values yield finer resolution at the cost of increased computation time. Defaults to 101, which is approximately equivalent to the default <code>incr = 0.01</code> used in <code>miso</code> .
<code>mu0</code>	a numeric value specifying the prior mean of the response. Defaults to 0.
<code>sig0</code>	a positive numeric value specifying the prior scale parameter, interpreted as the prior standard deviation of the response. Defaults to 100, yielding a diffuse prior.

kap0	a positive numeric value specifying the prior pseudo sample size for the mean. Smaller values yield a more diffuse prior on the mean. Defaults to 0.01.
nu0	a positive numeric value specifying the prior degrees of freedom for the variance. Smaller values yield a more diffuse prior on the variance. Defaults to 0.01.

Details

The prior distribution assumes that the mean response μ at each feature combination follows a Normal-Inverse-Chi-Squared model. Specifically, the conditional prior on μ given variance σ^2 is $\mu|\sigma^2 \sim N(\mu_0, \sigma^2/\kappa_0)$, and the marginal prior on μ is a t-distribution centered at μ_0 with ν_0 degrees of freedom and scale sig_0 . The default values $\text{kap}_0 = \nu_0 = 0.01$ approximate a Jeffreys non-informative prior, minimizing the influence of the prior on the posterior especially for feature combinations with few observations ($M = 1$). A more informative prior can be specified by increasing kap_0 and ν_0 .

Value

A list containing the following components:

- alldoses** a numeric matrix of unique feature combinations observed in the training data
- M** a numeric vector of observation counts at each feature combination
- MY** a numeric vector of sample means of the outcome at each feature combination (NA if $M = 0$)
- VY** a numeric vector of sample variances of the outcome at each feature combination (NA if $M \leq 1$)
- thetahat** a numeric vector of estimated mean responses at each unique feature combination, monotone nondecreasing with respect to the partial ordering of the features
- nt** an integer giving the number of threshold grid points used
- logH** a numeric vector of length nt giving the log posterior gain of the optimal classification at each threshold grid point

References

Cheung YK, Diaz KM. Monotone response surface of multi-factor condition: estimation and Bayes classifiers. *J R Stat Soc Series B Stat Methodol.* 2023 Apr;85(2):497-522. doi: 10.1093/jrsssb/qkad014. Epub 2023 Mar 22. PMID: 38464683; PMCID: PMC10919322.

Examples

```
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A), size=500, replace=TRUE),]
y <- rowSums(X) + rnorm(500)
misoN(X, y)
```

PBclassifier

*Multivariable Isotonic Classification using Projective Bayes***Description**

Estimates a monotone binary classification rule for multivariate features using a projective Bayes classifier. The classifier is obtained by projecting an unconstrained nonparametric Bayes estimator onto the partial ordering subspace defined by the assumption that the outcome probability is nondecreasing in each feature. The projection is computed using a recursive sequential update algorithm that yields the exact Bayes solution maximizing the posterior gain. Both a down-up ("DU") and an up-down ("UD") algorithm are available.

Usage

```
PBclassifier(X, y, method = "DU", a0 = 0.5, b0 = 0.5, t0 = 0.5)
```

Arguments

<code>X</code>	a numeric matrix of observed feature combinations, one row per observation, where repeated rows are expected. Each column represents a feature (e.g., a dose component or experimental factor) and each row represents the feature combination observed for one unit.
<code>y</code>	a binary numeric vector of length <code>nrow(X)</code> indicating the observed outcome for each observation (1 = event, 0 = no event).
<code>method</code>	a character string specifying the search strategy for finding the optimal monotone classification, either "DU" (down-up) or "UD" (up-down). Defaults to "DU".
<code>a0</code>	a positive numeric value specifying the shape1 hyperparameter of the Beta prior in the Beta-Binomial conjugate model. Defaults to 0.5 (Jeffreys prior).
<code>b0</code>	a positive numeric value specifying the shape2 hyperparameter of the Beta prior in the Beta-Binomial conjugate model. Defaults to 0.5 (Jeffreys prior).
<code>t0</code>	a numeric value in (0,1) specifying the threshold on the response rate used to classify each feature combination as event or no-event. Defaults to 0.5.

Value

A list of class "pbc" containing the following components:

alldoses	a numeric matrix of unique feature combinations
M	a numeric vector of observation counts at each feature combination
S	a numeric vector of event counts at each feature combination
yhat	a binary numeric vector of the optimal monotone classification for each feature combination (1 = event, 0 = no event)
pt	a numeric vector of posterior probabilities that the true response rate exceeds <code>t0</code> at each feature combination
logH	the log posterior probability of the optimal classification

References

Cheung YK, Diaz KM. Monotone response surface of multi-factor condition: estimation and Bayes classifiers. *J R Stat Soc Series B Stat Methodol.* 2023 Apr;85(2):497-522. doi: 10.1093/jrsssb/qkad014. Epub 2023 Mar 22. PMID: 38464683; PMCID: PMC10919322.

Cheung YK, Kuhn L. Evaluating multiplex diagnostic test using partially ordered Bayes classifier. *Ann Appl Stat.* In press.

Examples

```
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A),size=500, replace = TRUE),]
y <- as.numeric(rowSums(X)>=3)
PBclassifier(X,y)
```

predict.pbc

Predict Method for Projective Bayes Classifier

Description

Generates predictions from a fitted projective Bayes classifier for new feature combinations. Exact matches to training feature combinations are classified directly from the fitted classifier. For unobserved feature combinations, monotonicity constraints are used to impute the classification when the new combination is bounded from above or below by training combinations with a determinate classification. Combinations that cannot be classified by monotonicity constraints alone are flagged as indeterminate.

Usage

```
## S3 method for class 'pbc'
predict(object, Xnew, ...)
```

Arguments

object	a fitted object of class "pbc", produced by <code>PBclassifier</code> .
Xnew	a numeric matrix of new feature combinations to classify, with the same number of columns as the training data.
...	additional arguments (not used).

Value

A list containing the following components:

fit a numeric vector of predicted classifications (1 = event, 0 = no event, NA = indeterminate) for each row of Xnew

msg a character vector describing how each prediction was obtained: "exact match" for training combinations, "bound from below" or "bound from above" for combinations classified by monotonicity constraints, and "Unsure - need more training data" for indeterminate cases

Examples

```
A <- as.matrix(expand.grid(rep(list(0:1), 6)))
set.seed(2025)
X <- A[sample(nrow(A), size=500, replace=TRUE),]
y <- as.numeric(rowSums(X) >= 3)
fit <- PBclassifier(X, y)
predict(fit, X)
```

Index

boundary, [2](#)

mcmiso, [3](#), [5](#)

mcmisoN, [4](#)

mcPBclassifier, [2](#), [6](#)

miso, [3](#), [4](#), [8](#), [9](#)

misoN, [4](#), [5](#), [9](#)

PBclassifier, [2](#), [6](#), [7](#), [11](#), [12](#)

predict.pbc, [12](#)