

Package ‘CCI’

September 23, 2026

Type Package

Title Computational Test for Conditional Independence

Version 0.3.7

Date 2026-09-23

Maintainer Christian Thorjussen <christianbern@gmail.com>

Description Tool for performing computational testing for conditional independence between variables in a dataset. 'CCI' implements permutation in combination with Monte Carlo Cross-Validation in generating null distributions and test statistics. For more details see Computational Test for Conditional Independence (2024) <doi:10.3390/a17080323>.

Imports ggplot2, dplyr, caret, xgboost, ranger, stats, data.table, e1071, rlang, progress, kkn

Suggests testthat, knitr, rmarkdown, dagitty

License GPL (>= 2)

URL <https://github.com/khliland/CCI>

BugReports <https://github.com/khliland/CCI/issues>

Encoding UTF-8

LazyData true

RoxygenNote 7.3.3

Depends R (>= 3.5)

VignetteBuilder knitr

Config/testthat/edition 3

NeedsCompilation no

Author Christian Thorjussen [aut, cre] (ORCID:
<<https://orcid.org/0009-0005-5006-6491>>),
Kristian Hovde Liland [aut] (ORCID:
<<https://orcid.org/0000-0001-6468-9423>>)

Repository CRAN

Date/Publication 2026-09-23 13:00:02 UTC

Contents

add_interaction_terms	2
add_poly_terms	3
build_formula	4
CCI.direction	5
CCI.pretuner	7
CCI.test	10
check_formula	14
clean_formula	15
ExponentialNoise	15
get_pvalues	16
get_tuned_params	17
HardCase	17
is_categorical_Z_any	18
make_strata_from_categorical_Z	19
NonLinearCategorization	19
NonLinNormal	20
NonLinNormalZs_d0	21
NonLinNormalZs_d05	22
NormalData	23
perm.test	23
permute_within_strata	26
plot.CCI	26
PoissonNoise	28
PolyData	28
print.CCI	29
QQplot	30
test.gen	31
unclean_formula	34
UniformNoise_large	34
wrapper_knn	35
wrapper_ranger	36
wrapper_svm	37
wrapper_xgboost	38
Index	40

add_interaction_terms *Creates interaction terms for specified variables in a data frame Interaction terms are named as <var1>_int_<var2> (e.g., Z1_int_Z2 for the product of Z1 and Z2).*

Description

Creates interaction terms for specified variables in a data frame Interaction terms are named as <var1>_int_<var2> (e.g., Z1_int_Z2 for the product of Z1 and Z2).

Usage

```
add_interaction_terms(data, Z, mode = c("numeric_only", "mixed"))
```

Arguments

data	Data frame. The data frame containing the variables for which interaction terms are to be created.
Z	Character vector. The names of the variables for which interaction terms are to be created.
mode	Character. Specifies the type of interaction terms to create. Options are: <code>numeric_only</code> (only numeric-numeric interactions as products) or <code>mixed</code> (numeric-numeric as products, factor/character involved as categorical interactions). Default is <code>"numeric_only"</code> .

Value

A list with two components:

- `data`: The modified data frame with added interaction terms.
- `new_terms`: A character vector of the names of the added interaction terms (e.g., `Z1_int_2`).

Examples

```
data_generator <- function(N){
  Z1 <- rnorm(N,0,1)
  Z2 <- rnorm(N,0,1)
  X <- rnorm(N, Z1 + Z2, 1)
  Y <- rnorm(N, Z1 + Z2, 1)
  df <- data.frame(Z1, Z2, X, Y)
  return(df)
}
dat <- data_generator(250)
interaction_terms <- add_interaction_terms(data = dat, Z = c("Z1", "Z2"))
head(interaction_terms$data$Z1_int_Z2)
```

<code>add_poly_terms</code>	<i>Creates polynomial terms for specified variables in a data frame Polynomial terms are named as <variable>_d_<degree> (e.g., <code>Z1_d_2</code> for the square of <code>Z1</code>).</i>
-----------------------------	--

Description

Creates polynomial terms for specified variables in a data frame Polynomial terms are named as `<variable>_d_<degree>` (e.g., `Z1_d_2` for the square of `Z1`).

Usage

```
add_poly_terms(data, Z, degree = 3, poly = TRUE)
```

Arguments

data	Data frame. The data frame containing the variables for which polynomial terms are to be created.
Z	Character vector. The names of the variables for which polynomial terms are to be created. Only the numeric variables among them get polynomial terms; categorical variables (factors) are skipped.
degree	Integer. The maximum degree of polynomial terms to be created. Default is 3.
poly	Logical. If TRUE, polynomial terms will be created. If FALSE, no polynomial terms will be created. Default is TRUE.

Value

A list with two components:

- data: The modified data frame with added polynomial terms.
- new_terms: A character vector of the names of the added polynomial terms (e.g., Z1_d_2).

```
#'
```

Examples

```
set.seed(123)
data_generator <- function(N){
  Z1 <- rnorm(N,0,1)
  Z2 <- rnorm(N,0,1)
  X <- rnorm(N, Z1 + Z2, 1)
  Y <- rnorm(N, Z1 + Z2, 1)
  df <- data.frame(Z1, Z2, X, Y)
  return(df)
}
dat <- data_generator(250)
poly_terms <- add_poly_terms(data = dat, Z = c("Z1", "Z2"), degree = 3, poly = TRUE)
print(poly_terms$new_terms)
```

 build_formula

Build an expanded formula with poly and interaction terms

Description

Build an expanded formula with poly and interaction terms

Usage

```
build_formula(formula, poly_terms = NULL, interaction_terms = NULL)
```

Arguments

formula A base formula in the format $Y \sim X | Z1 + Z2$
poly_terms Character vector of polynomial term names
interaction_terms Character vector of interaction term names

Value

A formula object combining all terms

Examples

```
poly_terms <- c("Z1_d_2", "Z2_d_2")
interaction_terms <- c("Z1_int_Z2")
formula <- Y ~ X | Z1 + Z2
final_formula <- build_formula(formula, poly_terms, interaction_terms)
print(final_formula)
```

CCI.direction	<i>Choose Direction for testing for the CCI test</i>
---------------	--

Description

This function selects the best direction for the CCI test based on cross validation. For the condition $Y || X | Z$, the function return the recommended formula either $Y \sim X | Z$ or $X \sim Y | Z$.

Usage

```
CCI.direction(
  formula,
  data,
  method = "rf",
  folds = 4,
  nrounds = 600,
  max_depth = 6,
  eta = 0.3,
  gamma = 0,
  colsample_bytree = 1,
  min_child_weight = 1,
  MC_sample = 1,
  mtry = NULL,
  nthread = 1,
  k = 15,
  kernel = "optimal",
  distance = 2,
  poly = TRUE,
  degree = 3,
```

```

    interaction = TRUE,
    verbose = FALSE,
    subsample = NULL,
    ...
)

```

Arguments

formula	A formula of the form $Y \sim X \mid Z1 + Z2$ (or $Y \sim X + Z1 + Z2$, or $Y \sim X \mid 1$ without conditioning variables).
data	A data frame containing the variables specified in the formula.
method	Character. The learner: "rf" (random forest, ranger), "xgboost", "svm" (e1071) or "KNN" (kknn). Default is "rf".
folds	Integer. The number of folds for cross-validation. Default is 4.
nrounds	Integer. The number of trees for "rf" and boosting rounds for "xgboost". Default is 600.
max_depth	Integer. The maximum depth of the trees for xgboost. Default is 6.
eta	Numeric. The learning rate for xgboost. Default is 0.3.
gamma	Numeric. The minimum loss reduction required to make a further partition on a leaf node for xgboost. Default is 0.
colsample_bytree	Numeric. The subsample ratio of columns when constructing each tree for xgboost. Default is 1.
min_child_weight	Numeric. The minimum sum of instance weight (hessian) needed in a child for xgboost. Default is 1.
MC_sample	Numeric between 0 and 1. The share of the data to use (a random sample). Default is 1 (all data).
mtry	Integer. The number of variables tried at each split for "rf". Default is NULL (ranger's default).
nthread	Integer. The number of threads for "rf" and "xgboost". Default is 1.
k	Integer. The number of neighbours for "KNN". Default is 15.
kernel	Character. The kernel for "KNN". Default is "optimal".
distance	Numeric. The Minkowski distance parameter for "KNN". Default is 2.
poly	Not used. Kept for backward compatibility; CCI.test adds polynomial terms before choosing the direction.
degree	Not used. Kept for backward compatibility.
interaction	Not used. Kept for backward compatibility.
verbose	Logical. If TRUE, prints the cross-validated RMSE of both directions and the selected formula. Default is FALSE.
subsample	Deprecated, use MC_sample.
...	Additional arguments passed to the learner's wrapper function.

Details

Both directions are evaluated with k-fold cross-validation, using the same learners as `CCI.test` (`wrapper_ranger`, `wrapper_xgboost`, `wrapper_svm` or `wrapper_knn`) and the same folds for both directions. The numeric variables are standardized (mean 0, standard deviation 1) before the models are fitted, so the cross-validated RMSEs of the two directions are comparable and the choice does not depend on the units of the variables. The direction with the lowest standardized RMSE, i.e. the variable that is easiest to predict, is chosen.

Value

The selected formula, either $Y \sim X \mid Z$ or $X \sim Y \mid Z$. With `verbose = TRUE`, the cross-validated standardized RMSE of both directions is printed.

Examples

```
set.seed(1)
dat <- data.frame(Z1 = rnorm(200), Z2 = rnorm(200))
dat$X <- dat$Z1 + dat$Z2 + rnorm(200, sd = 0.1) # X is easy to predict from Z
dat$Y <- dat$Z1 + rnorm(200)                  # Y is harder to predict
CCI.direction(Y ~ X | Z1 + Z2, data = dat, method = "xgboost", nrounds = 50)
```

CCI.pretuner

CCI tuner function for CCI test

Description

The `CCI.pretuner` function performs a (random) grid search over hyperparameters for the machine learning models used by `CCI.test`. Each parameter combination is evaluated with the same model wrappers as used in `CCI.test` (`wrapper_ranger`, `wrapper_xgboost` and `wrapper_svm`), and all combinations are evaluated on the same resampling splits, so differences in performance reflect the parameters and not the random splits.

Usage

```
CCI.pretuner(
  formula,
  data,
  method = "rf",
  metric = "RMSE",
  validation_method = "cv",
  folds = 4,
  training_share = 0.7,
  tune_length = NULL,
  random_grid = TRUE,
  samples = 35,
  poly = TRUE,
  degree = 3,
```

```

interaction = TRUE,
verboseIter = NULL,
include_explanatory = FALSE,
verbose = FALSE,
progress = TRUE,
parallel = NULL,
nthread = 1,
num_trees = 500,
mtry = NULL,
nrounds = seq(100, 1000, by = 100),
eta = seq(0.01, 0.3, by = 0.05),
max_depth = 2:6,
gamma = c(0, 1, 2, 3),
colsample_bytree = c(0.8, 0.9, 1),
min_child_weight = c(1, 3),
subsample = NULL,
sigma = seq(0.1, 2, by = 0.3),
C = seq(0.1, 2, by = 0.5),
...
)

```

Arguments

formula	Model formula of the form $Y \sim X \mid Z1 + Z2$ (or $Y \sim X + Z1 + Z2$).
data	A data frame containing the variables specified in the formula.
method	Character. The machine learning method to tune: random forest "rf", extreme gradient boosting "xgboost" or Support Vector Machine "svm". Default is "rf".
metric	Character. The performance metric to optimize: "RMSE" (continuous outcome), "Kappa" or "LogLoss" (categorical outcome). Default is "RMSE".
validation_method	Character. The resampling method. One of "cv" (k-fold cross-validation), "LGOCV" (leave-group-out / Monte Carlo cross-validation) or "boot" (bootstrap, evaluated on the out-of-bag observations). Default is "cv".
folds	Integer. Number of folds ("cv") or resamples ("LGOCV", "boot"). Default is 4.
training_share	Numeric. Training share for validation_method = "LGOCV". Default is 0.7.
tune_length	Deprecated and ignored. Use samples to control the number of parameter combinations tried.
random_grid	Logical. If TRUE, samples combinations are drawn at random from the full grid. If FALSE, the full grid is evaluated. Default is TRUE.
samples	Integer. The number of random combinations to draw from the grid when random_grid = TRUE. Default is 35.
poly	Logical. If TRUE, polynomial terms of the conditioning variables are added. Set to FALSE if data already contains them (as when called from CCI.test). Default is TRUE.

degree	Integer. The degree of polynomial terms to include if poly is TRUE. Default is 3.
interaction	Logical. If TRUE, pairwise interaction terms of the conditioning variables are added. Default is TRUE.
verboseIter	Deprecated and ignored.
include_explanatory	Logical. If TRUE, given the condition $Y \parallel X \mid Z$, the explanatory variable X is included when tuning the model for Y. Default is FALSE.
verbose	Logical. If TRUE, prints details about the tuning process. Default is FALSE.
progress	Logical. If TRUE, shows a progress bar. Default is TRUE.
parallel	Deprecated and ignored. Use nthread for multithreaded model fitting.
nthread	Integer. Number of threads used when fitting rf and xgboost models. Default is 1.
num_trees	Integer. Number of trees used when tuning random forest. Default is 500.
mtry	Integer vector. Candidate values for the number of variables sampled at each split in random forest. Default is NULL, which uses $1:\min(p, \max(5, \text{ceiling}(p/2)))$ where p is the number of predictors.
nrounds	Integer vector. Candidate numbers of boosting rounds for xgboost. Default is <code>seq(100, 1000, by = 100)</code> .
eta	Numeric vector. Candidate learning rates for xgboost. Default is <code>seq(0.01, 0.3, by = 0.05)</code> .
max_depth	Integer vector. Candidate maximum tree depths for xgboost. Default is 2:6.
gamma	Numeric vector. Candidate minimum loss reductions for xgboost. Default is <code>c(0, 1, 2, 3)</code> .
colsample_bytree	Numeric vector. Candidate column subsample ratios for xgboost. Default is <code>c(0.8, 0.9, 1)</code> .
min_child_weight	Numeric vector. Candidate minimum child weights for xgboost. Default is <code>c(1, 3)</code> .
subsample	Deprecated and ignored.
sigma	Numeric vector. Candidate RBF kernel widths for svm (passed as gamma to <code>e1071::svm</code>). Default is <code>seq(0.1, 2, by = 0.3)</code> .
C	Numeric vector. Candidate cost (regularization) values for svm (passed as cost to <code>e1071::svm</code>). Default is <code>seq(0.1, 2, by = 0.5)</code> .
...	Additional arguments passed to the model wrapper. A custom grid can be supplied as <code>tuneGrid</code> , a data frame with one column per parameter (<code>mtry</code> for rf; <code>nrounds</code> , <code>eta</code> , <code>max_depth</code> , <code>gamma</code> , <code>colsample_bytree</code> , <code>min_child_weight</code> for xgboost; <code>sigma</code> , <code>C</code> for svm).

Details

By default the model is tuned for predicting Y from the conditioning set Z only (`include_explanatory = FALSE`), i.e. under the null hypothesis.

Value

A list containing:

- `best_param`: A one-row data frame with the best parameters, their mean performance and its standard deviation across resamples, and the method.
- `tuning_result`: A data frame with all tested parameter combinations and their performance, best first.
- `warnings`: A character vector of warnings and errors issued during tuning.

See Also

[CCI.test](#), [get_tuned_params](#), [perm.test](#)

Examples

```
set.seed(123)
data <- data.frame(x1 = rnorm(100), x2 = rnorm(100), x3 = rnorm(100), y = rnorm(100))
# Tune random forest parameters
result <- CCI.pretuner(formula = y ~ x1 | x2 + x3,
                      data = data,
                      samples = 5,
                      folds = 3,
                      method = "rf",
                      progress = FALSE)

result$best_param
```

CCI.test

Computational test for conditional independence based on ML and Monte Carlo Cross Validation

Description

The `CCI.test` function performs a conditional independence test using a specified machine learning model or a custom model provided by the user. It calculates the test statistic, generates a null distribution via permutations, computes p-values, and optionally generates a plot of the null distribution with the observed test statistic. The `'CCI.test'` function serves as a wrapper around the `'perm.test'` function

Usage

```
CCI.test(
  formula = NULL,
  data,
  p = 0.5,
  nperm = 160,
  nrounds = 600,
  mtry = NULL,
```

```

metric = "Auto",
method = "rf",
choose_direction = FALSE,
parametric = FALSE,
poly = TRUE,
degree = 3,
robust = TRUE,
MC_sample = "Auto",
MC_sample_set = NULL,
min_child_weight = 1,
colsample_bytree = 1,
eta = 0.3,
gamma = 0,
max_depth = 6,
interaction = TRUE,
mode = "numeric_only",
metricfunc = NULL,
mlfunc = NULL,
tail = NA,
tune = FALSE,
samples = 35,
folds = 5,
tune_length = NULL,
k = 15,
center = TRUE,
scale = TRUE,
eps = 1e-15,
positive = NULL,
kernel = "optimal",
distance = 2,
seed = NA,
random_grid = TRUE,
nthread = 2,
verbose = FALSE,
progress = TRUE,
subsample = NULL,
subsample_set = NULL,
...
)

```

Arguments

formula	Model formula specifying the relationship between dependent and independent variables. (Ex: $Y \sim X \mid Z1 + Z2$ for $Y \parallel X \mid Z1, Z2$). For an unconditional test of $Y \parallel X$, write $Y \sim X \mid 1$ or $Y \sim X + 1$; $Y \sim X$ alone gives an error.
data	A data frame containing the variables specified in the formula.
p	Numeric. Proportion of data used for training the model. Default is 0.5.
nperm	Integer. The number of permutations to perform. Default is 60.

nrounds	Integer. The number of rounds (trees) for methods 'xgboost' and 'rf'. Default is 600.
mtry	Number of variables to possibly split at in each node for method 'rf'. Default is NULL (sqrt of number of variables).
metric	Character. Specifies the type of data: "Auto", "RMSE" or "Kappa". Default is "Auto".
method	Character. Specifies the machine learning method to use. Supported methods are random forest "rf", extreme gradient boosting "xgboost", support vector machine 'svm' and K-nearest neighbour 'KNN'. Default is "rf".
choose_direction	Logical. If TRUE, the function will choose the best direction for testing. Default is FALSE.
parametric	Logical, indicating whether to compute a parametric p-value instead of the empirical p-value. A parametric p-value assumes that the null distribution is gaussian. Default is FALSE.
poly	Logical. If TRUE, polynomial terms of the conditional variables are included in the model. Default is TRUE.
degree	Integer. The degree of polynomial terms to include if poly is TRUE. Default is 3.
robust	Logical. If TRUE and the conditioning set Z contains any categorical variables (factor, character or logical), X is permuted within the groups defined by the categorical variables in Z (stratified permutation). This keeps the relationship between X and the categorical part of Z under the null hypothesis. If FALSE, X is always permuted over all observations. Default is TRUE.
MC_sample	Character. The share of the data used in each Monte Carlo sample: "Auto" uses $(900/n)^{0.75}$ of the data when $n > 900$ and all data otherwise, "Yes" uses the share given in MC_sample_set, and "No" always uses all data. A different random sample is drawn in each Monte Carlo sample. Default is "Auto".
MC_sample_set	Numeric between 0 and 1. The share of the data used in each Monte Carlo sample when MC_sample = "Yes".
min_child_weight	Numeric. The minimum sum of instance weight (hessian) needed in a child for methods like xgboost. Default is 1.
colsample_bytree	Numeric. The subsample ratio of columns when constructing each tree for methods like xgboost. Default is 1.
eta	Numeric. The learning rate for methods like xgboost. Default is 0.3.
gamma	Numeric. The minimum loss reduction required to make a further partition on a leaf node of the tree for methods like xgboost. Default is 0.
max_depth	Integer. The maximum depth of the trees for methods like xgboost. Default is 6.
interaction	Logical. If TRUE, interaction terms of the conditional variables are included in the model. Default is TRUE.
mode	Character. Specifies the mode of operation: "numeric_only" or "mixed". Default is "numeric_only".

metricfunc	Optional custom performance metric: a function function(actual, predictions, ...) returning a single number. Set tail to "right" if higher values mean better predictions and "left" if lower values do. actual is numeric for a numeric Y and a factor for a categorical Y. For a numeric Y, predictions is numeric. For a categorical Y, predictions are the predicted classes (a factor) for methods "rf", "svm" and "KNN", and class probabilities for "xgboost": the probability of the second class level for two classes, or an n x K matrix with the class levels as column names for more classes. The ... arguments are only passed on if the function accepts them. Default is NULL.
mlfunc	Optional the user can pass a custom machine learning wrapper function to use instead of the predefined methods. Default is NULL.
tail	Character. Specifies whether to calculate left-tailed or right-tailed p-values, depending on the performance metric used. Only applicable if using metricfunc or mlfunc. Default is NA.
tune	Logical. If TRUE, hyperparameters for the specified machine learning method are tuned with CCI.pretuner before testing, and the best parameters are used in the test. Available for methods 'rf', 'xgboost' and 'svm'. Default is FALSE.
samples	Integer. Number of hyperparameter combinations used in tuning. Default is 35.
folds	Integer. The number of folds for cross-validation during the tuning process. Default is 5.
tune_length	Deprecated and ignored. Use samples to control the number of parameter combinations tried in tuning.
k	Integer. The number of nearest neighbors to use for KNN method. Default is 15.
center	Logical. If TRUE, the data will be centered before fitting the model
scale	Logical. If TRUE, the data will be scaled before fitting the model. Default is TRUE.
eps	Numeric. A small value to avoid division by zero in some calculations.
positive	Character. The name of the positive class (KNN) in the data, used for classification tasks. Default is NULL.
kernel	Character. The kernel type to use for KNN method. Default is "optimal".
distance	Numeric. Parameter of Minkowski distance for the "KNN" method. Default is 2.
seed	Integer. Set the seed for reproducing results. Default is NA.
random_grid	Logical. If TRUE, a random grid search is performed. If FALSE, a full grid search is performed. Default is TRUE.
nthread	Integer. The number of threads to use for parallel processing. Default is 1.
verbose	Logical. If TRUE, additional information is printed during the execution of the function. Default is FALSE.
progress	Logical. If TRUE, a progress bar is displayed during the permutation process. Default is TRUE.
subsample	Deprecated, use MC_sample.
subsample_set	Deprecated, use MC_sample_set.
...	Additional arguments to pass to the perm.test function.

Value

Invisibly returns the result of `perm.test`, which is an object of class 'CCI' containing the null distribution, observed test statistic, p-values, the machine learning model used, and the data.

See Also

[perm.test](#), [print.summary.CCI](#), [plot.CCI](#), [CCI.pretuner](#), [QQplot](#)

Examples

```
set.seed(123)
data <- data.frame(x1 = stats::rnorm(100), x2 = stats::rnorm(100), y = stats::rnorm(100))
result <- CCI.test(y ~ x1 | x2, data = data, nperm = 25, interaction = FALSE)
summary(result)
```

check_formula

Check the formula statement

Description

This function verifies that all variables specified in the formula are present in the provided data frame. If any variables are missing, the function will stop and return an error message listing the missing variables.

Usage

```
check_formula(formula, data)
```

Arguments

formula	Formula. The model formula that specifies the relationship between the dependent and independent variables.
data	Data frame. The data frame in which to check for the presence of variables specified in the formula.

Value

Invisibly returns NULL if all variables are present. Stops with an error if any variables are missing.

clean_formula	<i>Clean and Reformat Formula String</i>
---------------	--

Description

This function processes and reformats formula string to ensure it is in the correct format for conditional independence testing. The function checks if the formula uses the '+' operator for additive models and transforms it into a format that includes a conditioning variable separated by '|'.

Usage

```
clean_formula(formula)
```

Arguments

formula	Formula. The model formula that specifies the relationship between the dependent and independent variables, and potentially the conditioning variables. The formula is expected to follow the format $Y \sim X + Z1 + Z2$ or $Y \sim X Z1 + Z2$, or $Y \sim X 1$ or $Y \sim X + 1$ for an unconditional test.
---------	--

Details

For an unconditional test of $Y \parallel X$ (no conditioning variables), the formula must be written $Y \sim X | 1$ or $Y \sim X + 1$; both are returned as $Y \sim X | 1$. A formula without conditioning variables and without the explicit 1, such as $Y \sim X$, gives an error.

Value

A reformatted formula in the correct format for conditional independence testing. The returned formula will either retain the original format or be transformed to include conditioning variables.

Examples

```
clean_formula(y ~ x | z + v)
clean_formula(y ~ x + z + v)
clean_formula(y ~ x + 1) # unconditional test
```

ExponentialNoise	<i>Example dataset: ExponentialNoise</i>
------------------	--

Description

A dataset containing simulated conditional independence test results.

Usage

```
ExponentialNoise
```

Format

A data frame with 600 rows and 3 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

Source

Simulated data.

get_pvalues

P-value Calculation Based on Null Distribution and Test Statistic

Description

This function calculates p-values based on the comparison of a test statistic against a null distribution. It can perform either empirical or parametric p-value calculations and supports both left-tailed and right-tailed tests.

Usage

```
get_pvalues(
  dist,
  test_statistic,
  parametric = FALSE,
  tail = c("left", "right")
)
```

Arguments

dist	Numeric vector. Represents the null distribution of the test statistic.
test_statistic	Numeric. The observed test statistic for which the p-value is to be calculated.
parametric	Logical. If TRUE, calculates parametric p-values assuming the null distribution is normal. If FALSE, calculates empirical p-values. Default is FALSE.
tail	Character. Specifies whether to calculate left-tailed or right-tailed p-values. Must be either "left" or "right". Default is "left".

Details

Missing values (NA/NaN) in **dist**, typically from model fits that failed during the permutations, are removed with a warning, and the p-value is computed from the remaining values. If the test statistic is missing, or fewer than two valid null values remain, NA is returned with a warning.

Value

Numeric. The calculated p-value.

Examples

```
set.seed(123)
null_dist <- rnorm(1000)
observed_stat <- 1.5
p_value <- get_pvalues(null_dist, observed_stat, parametric = FALSE, tail = "right")
print(p_value)
```

get_tuned_params	<i>Get the best parameters after tuning with CCI.tuner</i>
------------------	--

Description

Get the best parameters after tuning with CCI.tuner

Usage

```
get_tuned_params(tuned_model)
```

Arguments

tuned_model	A model object returned from the CCI.pretuner function. This object contains the tuned parameters and other relevant information.
-------------	---

Value

A named list of tuned parameters specific to the model method (e.g., mtry for random forest, eta, max_depth for xgboost), named as the arguments of [CCI.test](#). Returns NULL for unsupported methods. The xgboost row subsample is not returned, since CCI.test has a deprecated argument subsample (the old name of MC_sample).

HardCase	<i>Example dataset: HardCase</i>
----------	----------------------------------

Description

A dataset containing simulated conditional independence test results.

Usage

```
HardCase
```

Format

A data frame with 500 rows and 3 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

Source

Simulated data.

is_categorical_Z_any *Check whether Z contains at least one categorical variable*

Description

Categorical is defined as factor (and optionally character and logical). This helper is intentionally strict about inputs and does not modify data.

Usage

```
is_categorical_Z_any(sub_data, Z, allow_character = TRUE, allow_logical = TRUE)
```

Arguments

sub_data	data.frame containing the Z columns.
Z	character vector of column names defining the conditioning set.
allow_character	logical; treat character as categorical. Default TRUE.
allow_logical	logical; treat logical as categorical. Default TRUE.

Value

logical scalar.

make_strata_from_categorical_Z

Create strata from the categorical subset of Z

Description

Uses `interaction()` on the categorical Z columns only. Characters (and logicals, if enabled) are coerced to factor for stable behavior.

Usage

```
make_strata_from_categorical_Z(
  sub_data,
  Z,
  allow_character = TRUE,
  allow_logical = TRUE
)
```

Arguments

<code>sub_data</code>	data.frame containing Z columns.
<code>Z</code>	character vector of Z column names.
<code>allow_character</code>	logical; treat character as categorical. Default TRUE.
<code>allow_logical</code>	logical; treat logical as categorical. Default TRUE.

Details

Rows with NA in any stratification column will receive NA strata.

Value

A factor defining strata (same length as `nrow(sub_data)`).

NonLinearCategorization

Example dataset: NonLinearCategorization

Description

A dataset containing simulated data from a non-linear transformation followed by categorization.

Usage

```
NonLinearCategorization
```

Format

A data frame with 600 rows and 3 variables:

X Numeric vector

Y Numeric vector

Z Conditioning variable

Source

Simulated data.

NonLinNormal

Example dataset: NonLinNormal

Description

A dataset containing simulated data from a non-linear transformation of a multivariate normal distribution.

Usage

NonLinNormal

Format

A data frame with 500 rows and 4 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

Source

Simulated data.

NonLinNormalZs_d0*Example dataset: NonLinNormalZs_d0*

Description

A dataset containing simulated data with uniform noise.

Usage

NonLinNormalZs_d0

Format

A data frame with 1000 rows and 15 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

Z3 Conditioning variable

Z4 Conditioning variable

Z5 Conditioning variable

Z6 Conditioning variable

Z7 Conditioning variable

Z8 Conditioning variable

Z9 Conditioning variable

Z10 Conditioning variable

Z11 Conditioning variable

Z12 Conditioning variable

Z13 Conditioning variable

Z14 Conditioning variable

Z15 Conditioning variable

Source

Simulated data.

NonLinNormalZs_d05	<i>Example dataset: NonLinNormalZs_d05</i>
--------------------	--

Description

A dataset containing simulated data with uniform noise.

Usage

NonLinNormalZs_d05

Format

A data frame with 1000 rows and 15 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

Z3 Conditioning variable

Z4 Conditioning variable

Z5 Conditioning variable

Z6 Conditioning variable

Z7 Conditioning variable

Z8 Conditioning variable

Z9 Conditioning variable

Z10 Conditioning variable

Z11 Conditioning variable

Z12 Conditioning variable

Z13 Conditioning variable

Z14 Conditioning variable

Z15 Conditioning variable

Source

Simulated data.

NormalData*Example dataset: NormalData*

Description

A dataset containing simulated data from a multivariate normal distribution.

Usage

NormalData

Format

A data frame with 400 rows and 4 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

@source Simulated data.

perm.test*Permutation Test for Conditional Independence*

Description

Permutation Test for Conditional Independence

Usage

```
perm.test(  
  formula,  
  data,  
  p = 0.5,  
  nperm = 160,  
  MC_sample = 1,  
  metric = "RMSE",  
  method = "rf",  
  nrounds = 600,  
  mtry = NULL,  
  parametric = FALSE,  
  tail = NA,  
  robust = TRUE,  
  metricfunc = NULL,
```

```

mlfunc = NULL,
nthread = 1,
progress = TRUE,
k = 15,
center = TRUE,
scale = TRUE,
eps = 1e-15,
positive = NULL,
kernel = "optimal",
distance = 2,
subsample = NULL,
...
)

```

Arguments

formula	Model formula or DAGitty object specifying the relationship between dependent and independent variables.
data	A data frame containing the variables specified in the formula.
p	Proportion of data to use for training the model. Default is 0.5.
nperm	Number of permutations to perform. Default is 160.
MC_sample	Numeric between 0 and 1. The share of the data used in each Monte Carlo sample (a new random sample each time). Default is 1 (all data).
metric	Type of metric: "RMSE", "Kappa" or "LogLoss". Default is 'RMSE'.
method	The machine learning method to use for the learner. Supported methods include "rf", "xgboost", "KNN" and "svm". Default is "rf".
nrounds	Number of rounds (trees) for methods 'xgboost' and 'rf'. Default is 600.
mtry	Number of variables to possibly split at in each node for method 'rf'. Default is NULL (sqrt of number of variables).
parametric	Logical. If TRUE, a parametric p-value is calculated instead of an empirical p-value. Default is FALSE.
tail	Specifies whether the test is one-tailed ("left" or "right") or two-tailed. Default is NA.
robust	Logical. If TRUE and the conditioning set Z contains any categorical variables (factor, character or logical), X is permuted within the groups defined by the categorical variables in Z (stratified permutation). This keeps the relationship between X and the categorical part of Z under the null hypothesis. If FALSE, X is always permuted over all observations. Default is TRUE.
metricfunc	Optional custom performance metric: a function function(actual, predictions, ...) returning a single number. Set tail to "right" if higher values mean better predictions and "left" if lower values do. actual is numeric for a numeric Y and a factor for a categorical Y. For a numeric Y, predictions is numeric. For a categorical Y, predictions are the predicted classes (a factor) for methods "rf", "svm" and "KNN", and class probabilities for "xgboost": the probability of the second class level for two classes, or an n x K matrix with the class levels

	as column names for more classes. The ... arguments are only passed on if the function accepts them. Default is NULL.
mlfunc	An optional custom machine learning function to use instead of the predefined methods. Default is NULL.
nthread	Integer. The number of threads to use for parallel processing for method 'rf' and 'xgboost'. Default is 1.
progress	Logical. If TRUE, a progress bar is displayed during the permutation process. Default is TRUE.
k	Integer. The number of nearest neighbors for the "KNN" method. Default is 15.
center	Logical. If TRUE, the data is centered before model fitting. Default is TRUE.
scale	Logical. If TRUE, the data is scaled before model fitting. Default is TRUE.
eps	Numeric. A small value added to avoid division by zero. Default is 1e-15.
positive	Character vector. Specifies which levels of a factor variable should be treated as positive class in classification tasks. Default is NULL.
kernel	Character string specifying the kernel type for method option "KNN" . Possible choices are "rectangular" (which is standard unweighted knn), "triangular", "epanechnikov" (or beta(2,2)), "biweight" (or beta(3,3)), "triweight" (or beta(4,4)), "cos", "inv", "gaussian" and "optimal". Default is "optimal".
distance	Numeric. Parameter of Minkowski distance for the "KNN" method. Default is 2.
subsample	Deprecated, use MC_sample.
...	Additional arguments to pass to the machine learning model fitting function.

Value

An object of class 'CCI' containing the null distribution, observed test statistic, p-values, the machine learning model used, and the data. The element settings holds all arguments passed to [test.gen](#) (including custom metricfunc/mlfunc and model parameters), so that [QQplot](#) can repeat the test with the same settings.

See Also

[print.CCI](#), [summary.CCI](#), [plot.CCI](#), [QQplot](#)

Examples

```
set.seed(123)
dat <- data.frame(x1 = rnorm(100),
  x2 = rnorm(100),
  x3 = rnorm(100),
  x4 = rnorm(100),
  y = rnorm(100))
perm.test(y ~ x1 | x2 + x3 + x4, data = dat, nperm = 25)
```

`permute_within_strata` *Stratified permutation of x within strata*

Description

Permutes values of 'x' within each stratum. By default, rows with NA in 'x' or 'strata' are left unchanged to preserve alignment.

Usage

```
permute_within_strata(x, strata, seed = NULL, na_action = c("keep", "drop"))
```

Arguments

<code>x</code>	vector to permute.
<code>strata</code>	factor-like vector defining strata.
<code>seed</code>	optional seed for reproducibility.
<code>na_action</code>	how to handle NA rows: "keep" (default) keeps them fixed, "drop" removes them (returns shorter vector).

Value

`x_star` permuted within strata.

Examples

```
set.seed(123)
x <- 1:10
strata <- rep(letters[1:2], each = 5)
x_permuted <- permute_within_strata(x, strata, seed = 123)
```

`plot.CCI` *Plot for CCI testing*

Description

Plot for CCI testing

Usage

```
## S3 method for class 'CCI'
plot(
  x,
  ...,
  fill_color = "lightblue",
  title.size = 14,
  axis.text.x = 13,
  axis.text.y = 13,
  strip.text.x = 13,
  strip.text.y = 13,
  legend.text = 13,
  legend.title = 13,
  axis.title.x = 13,
  axis.title.y = 13,
  base_size = 13
)
```

Arguments

<code>x</code>	Object of class 'CCI'
<code>...</code>	ggplot2 layers or themes to add to the plot, e.g. <code>plot(result, ggplot2::labs(title = "My test"))</code> . Other arguments are ignored with a warning. Since the plot is a ggplot2 object, layers can also be added with <code>+</code> .
<code>fill_color</code>	Color for the histogram fill
<code>title.size</code>	Size of the plot title
<code>axis.text.x</code>	Size of x-axis text
<code>axis.text.y</code>	Size of y-axis text
<code>strip.text.x</code>	Size of x-axis strip text
<code>strip.text.y</code>	Size of y-axis strip text
<code>legend.text</code>	Size of legend text
<code>legend.title</code>	Size of legend title
<code>axis.title.x</code>	Size of x-axis title
<code>axis.title.y</code>	Size of y-axis title
<code>base_size</code>	Base font size

Value

A plot of the null distribution and the test statistic in ggplot2 format.

See Also

[print.CCI](#), [summary.CCI](#), [plot.CCI](#), [perm.test](#)

Examples

```
dat <- data.frame(x1 = rnorm(100), x2 = rnorm(100), y = rnorm(100))
cci <- CCI.test(y ~ x1 + x2, data = dat, interaction = FALSE, nperm = 30)
plot(cci)
```

PoissonNoise	<i>Example dataset: PoissonNoise</i>
--------------	--------------------------------------

Description

A dataset containing simulated data from a Poisson distribution.

Usage

PoissonNoise

Format

A data frame with 1000 rows and 4 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

Source

Simulated data.

PolyData	<i>Example dataset: PolyData</i>
----------	----------------------------------

Description

A dataset containing simulated data from a polynomial relationship.

Usage

PolyData

Format

A data frame with 600 rows and 4 variables:

X Numeric vector

Y Numeric vector

Z1 Conditioning variable

Z2 Conditioning variable

Source

Simulated data.

print.CCI	<i>Print and summary methods for the CCI class</i>
-----------	--

Description

print() gives a short overview of a CCI test: the hypothesis tested, the learner and metric, the test statistic and the p-value. summary() gives the same with the number of permutations, the tail of the test and the share of data used in each Monte Carlo sample (MC_sample) on separate lines.

When choose_direction = TRUE changed the direction of the test (e.g. from $Y \sim X \mid Z$ to $X \sim Y \mid Z$), the formula that was tested is shown, together with the formula as given.

Usage

```
## S3 method for class 'CCI'
print(x, digits = 4, ...)

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

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

Arguments

x	Object of class 'CCI' (for print) or 'summary.CCI' (for print.summary.CCI)
digits	Number of significant digits for the test statistic and p-value. Default is 4.
...	Additional arguments to print/summary
object	Object of class 'CCI'

Value

The print methods return their argument invisibly, the summary method returns an object of class 'summary.CCI'.

See Also

[perm.test](#), [plot.CCI](#), [QQplot](#)

Examples

```
set.seed(1)
dat <- data.frame(Z = rnorm(100), X = rnorm(100))
dat$Y <- dat$Z + rnorm(100)
res <- CCI.test(Y ~ X | Z, data = dat, nperm = 20, progress = FALSE)
res
summary(res)
```

 QQplot

QQ-plot for multiple testing in CCI

Description

QQ-plot for multiple testing in CCI

Usage

```
QQplot(
  object,
  title.size = 14,
  axis.text.x = 13,
  axis.text.y = 13,
  strip.text.x = 13,
  strip.text.y = 13,
  legend.text = 13,
  legend.title = 13,
  axis.title.x = 13,
  axis.title.y = 13,
  progress = TRUE,
  ...
)
```

Arguments

<code>object</code>	Object of class 'CCI'
<code>title.size</code>	Size of the plot title
<code>axis.text.x</code>	Size of x-axis text
<code>axis.text.y</code>	Size of y-axis text
<code>strip.text.x</code>	Size of x-axis strip text
<code>strip.text.y</code>	Size of y-axis strip text
<code>legend.text</code>	Size of legend text

legend.title	Size of legend title
axis.title.x	Size of x-axis title
axis.title.y	Size of y-axis title
progress	Logical indicating whether to show progress during computation
...	Arguments passed to test.gen . By default the test statistic is recomputed <code>nperm</code> times with exactly the settings used to create object (method, metric, custom <code>metricfunc/mlfunc</code> , model parameters, scaling, etc.). Arguments given here override them, e.g. <code>nperm = 50</code> .

Details

The test statistic (without permutation) is recomputed on `nperm` new random train/test splits, and the p-value of each is computed against the stored null distribution. Under the null hypothesis the p-values should be roughly uniform.

Value

A QQ-plot of the p-values in `ggplot2` format.

See Also

[print.CCI](#), [summary.CCI](#), [plot.CCI](#), [perm.test](#)

Examples

```
dat <- data.frame(x1 = rnorm(100), x2 = rnorm(100), y = rnorm(100))
cci <- CCI.test(y ~ x1 | x2,
  data = dat,
  nperm = 25,
  interaction = FALSE)
QQplot(cci)
```

test.gen

Generate the Test Statistic or Null Distribution Using Permutation

Description

This function generates the test statistic or a null distribution through permutation for conditional independence testing. It supports various machine learning methods, including random forests, extreme gradient boosting, and allows for custom metric functions and model fitting functions.

Usage

```
test.gen(
  formula,
  data,
  method = "rf",
  metric = "RMSE",
  nperm = 160,
  MC_sample = 1,
  p = 0.5,
  nrounds = 600,
  mtry = NULL,
  nthread = 1,
  permutation = FALSE,
  robust = TRUE,
  metricfunc = NULL,
  mlfunc = NULL,
  progress = TRUE,
  center = TRUE,
  scale = TRUE,
  eps = 1e-15,
  k = 15,
  positive = NULL,
  kernel = "optimal",
  distance = 2,
  subsample = NULL,
  ...
)
```

Arguments

formula	Formula specifying the relationship between dependent and independent variables.
data	Data frame. The data containing the variables used.
method	Character. The modeling method to be used. Options include "xgboost" for gradient boosting, or "rf" for random forests or "svm" for Support Vector Machine.
metric	Character. The type of metric: can be "RMSE", "Kappa" or "LogLoss". Default is 'RMSE'
nperm	Integer. The number of generated Monte Carlo samples. Default is 160.
MC_sample	Numeric between 0 and 1. The share of the data used in each Monte Carlo sample (a new random sample each time). Default is 1 (all data).
p	Numeric. The proportion of the data to be used for training. The remaining data will be used for testing. Default is 0.5.
nrounds	Integer. The number of rounds (trees) for methods like 'xgboost' and 'rf'. Default is 600.
mtry	Integer. The number of variables to possibly split at in each node for method 'rf'. Default is the rounded down square root of numbers of columns in data.

nthread	Integer. The number of threads to use for parallel processing. Only relevant for methods 'rf' and 'xgboost'. Default is 1.
permutation	Logical. Whether to perform permutation of the 'X' variable. Used to generate a null distribution. Default is FALSE.
robust	Logical. If TRUE and the conditioning set Z contains any categorical variables (factor, character or logical), X is permuted within the groups defined by the categorical variables in Z (stratified permutation). If FALSE, X is always permuted over all observations. Default is TRUE.
metricfunc	Optional custom performance metric: a function function(actual, predictions, ...) returning a single number. Set tail to "right" if higher values mean better predictions and "left" if lower values do. actual is numeric for a numeric Y and a factor for a categorical Y. For a numeric Y, predictions is numeric. For a categorical Y, predictions are the predicted classes (a factor) for methods "rf", "svm" and "KNN", and class probabilities for "xgboost": the probability of the second class level for two classes, or an n x K matrix with the class levels as column names for more classes. The ... arguments are only passed on if the function accepts them. Default is NULL.
mlfunc	Function. A custom machine learning function provided by the user. The function must have the arguments: formula, data, train_indices, test_indices, and ..., and return a single value performance metric. Default is NULL.
progress	Logical. A logical value indicating whether to show a progress bar during when building the null distribution. Default is TRUE.
center	Logical. If TRUE, the data is centered before model fitting. Default is TRUE.
scale	Logical. If TRUE, the data is scaled before model fitting. Default is TRUE.
eps	Numeric. A small value added to avoid division by zero. Only relevant for method 'KNN'. Default is 1e-15.
k	Integer. The number of nearest neighbors for the "KNN" method. Default is 15.
positive	Character vector. Only relevant for method 'KNN'. Specifies which levels of a factor variable should be treated as positive class in classification tasks. Default is NULL.
kernel	Character. Only relevant for method 'KNN'. Specifies the kernel type for method option "KNN". Possible choices are "rectangular" (which is standard unweighted knn), "triangular", "epanechnikov" (or beta(2,2)), "biweight" (or beta(3,3)), "triweight" (or beta(4,4)), "cos", "inv", "gaussian" and "optimal". Default is "optimal".
distance	Numeric. Parameter of Minkowski distance for the "KNN" method. Default is 2.
subsample	Deprecated, use MC_sample.
...	Additional arguments to pass to the machine learning wrapper functions wrapper_xgboost, wrapper_ranger, wrapper_knn and wrapper_svm, or to a custom-built wrapper function.

Value

A list containing the test distribution.

Examples

```

set.seed(123)
data <- data.frame(x1 = rnorm(100),
  x2 = rnorm(100),
  x3 = rnorm(100),
  x4 = rnorm(100),
  y = rnorm(100))
result <- test.gen(formula = y ~ x1 | x2 + x3 + x4,
  metric = "RMSE",
  nperm = 50,
  data = data)
hist(result$distribution)

```

unclean_formula	<i>Convert CI-style formula $Y \sim X \mid Z$ into regression-style $Y \sim X + Z$</i>
-----------------	--

Description

Convert CI-style formula $Y \sim X \mid Z$ into regression-style $Y \sim X + Z$

Usage

```
unclean_formula(f)
```

Arguments

f A formula of the form $Y \sim X \mid Z1 + Z2$, or already $Y \sim X + Z1 + Z2$.

Value

A standard formula of the form $Y \sim X + Z1 + Z2$.

UniformNoise_large	<i>Example dataset: UniformNoise_large</i>
--------------------	--

Description

A dataset containing simulated data with uniform noise.

Usage

```
UniformNoise_large
```

Format

- A data frame with 20000 rows and 4 variables:
- X** Numeric vector
- Y** Numeric vector
- Z1** Conditioning variable
- Z2** Conditioning variable

Source

Simulated data.

wrapper_knn	<i>k-Nearest Neighbors (KNN) wrapper for CCI (kkn-based)</i>
-------------	--

Description

k-Nearest Neighbors (KNN) wrapper for CCI (kkn-based)

Usage

```
wrapper_knn(  
  formula,  
  data,  
  train_indices,  
  test_indices,  
  metric,  
  metricfunc = NULL,  
  k = 15,  
  eps = 1e-15,  
  positive = NULL,  
  kernel = "optimal",  
  distance = 2,  
  ...  
)
```

Arguments

- | | |
|---------------|---|
| formula | Model formula |
| data | Data frame |
| train_indices | Indices for training rows |
| test_indices | Indices for test rows |
| metric | Performance metric: "RMSE" (regression), "Kappa" (classification), or "LogLoss" (classification), or the name of a custom metric when metricfunc is given |

metricfunc	Optional custom metric function function(actual, predictions, ...). A numeric response gives regression (numeric predictions); a factor, character or logical response gives classification (predicted classes as a factor). The ... arguments are only passed on if the function accepts them.
k	Integer, number of neighbors (default 15)
eps	Small value to avoid log(0) in LogLoss calculations. Default is 1e-15.
positive	Character. The positive class label for binary classification (used in LogLoss). Default is NULL.
kernel	Character. Weighting kernel for kkn. Default "optimal".
distance	Numeric. Minkowski distance parameter. 2 = Euclidean. Default 2.
...	Additional arguments passed to kkn::kkn (e.g., ykernel, na.action)

Value

Numeric performance metric

wrapper_ranger	<i>Random Forest wrapper for CCI</i>
----------------	--------------------------------------

Description

Random Forest wrapper for CCI

Usage

```

wrapper_ranger(
  formula,
  data,
  train_indices,
  test_indices,
  metric,
  metricfunc = NULL,
  nthread = 1,
  mtry = NULL,
  num.trees,
  eps = 1e-15,
  ...
)

```

Arguments

formula	Model formula specifying the dependent and independent variables.
data	Data frame containing the dataset to be used for training and testing the model.
train_indices	A vector of indices specifying the rows in data to be used as the training set.
test_indices	A vector of indices specifying the rows in data to be used as the test set.

metric	Type of metric ("RMSE", "Kappa" or "Log Loss")
metricfunc	Optional user-defined function function(actual, predictions, ...) returning a numeric performance value. For a factor response, predictions are the predicted classes. The ... arguments are only passed on if the function accepts them.
nthread	Integer. The number of threads to use for parallel processing. Default is 1.
mtry	Integer. The number of variables to possibly split at in each node. Default is the square root of the number of columns in data.
num.trees	Integer. The number of trees to grow in the random forest.
eps	Small value to avoid log(0) in LogLoss calculations. Default is 1e-15.
...	Additional arguments passed to the ranger function.

Value

A numeric value representing the performance metric of the model on the test set.

wrapper_svm	<i>SVM wrapper for CCI</i>
-------------	----------------------------

Description

SVM wrapper for CCI

Usage

```

wrapper_svm(
  formula,
  data,
  train_indices,
  test_indices,
  metric,
  metricfunc = NULL,
  eps = 1e-15,
  ...
)

```

Arguments

formula	Model formula
data	Data frame
train_indices	Indices for training data
test_indices	Indices for testing data
metric	Type of metric ("RMSE", "Kappa" or "Log Loss")
metricfunc	Optional user-defined function to calculate a custom performance metric.
eps	Small value to avoid log(0) in LogLoss calculations. Default is 1e-15.
...	Additional arguments passed to e1071::svm

Value

Performance metric (RMSE for continuous, Kappa for classification)

wrapper_xgboost	<i>Extreme Gradient Boosting wrapper for CCI</i>
-----------------	--

Description

Fits an xgboost model on the training rows and returns its performance on the test rows.

Usage

```
wrapper_xgboost(  
  formula,  
  data,  
  train_indices,  
  test_indices,  
  metric,  
  nrounds = 500,  
  metricfunc = NULL,  
  nthread = 1,  
  eps = 1e-15,  
  MC_sample = 1,  
  ...  
)
```

Arguments

formula	Model formula
data	Data frame
train_indices	Indices for training data
test_indices	Indices for test data
metric	Type of metric ("RMSE", "Kappa" or "LogLoss"), or the name of a custom metric when metricfunc is given
nrounds	Number of boosting rounds
metricfunc	Optional user-defined function function(actual, predictions, ...) returning a numeric performance value. See Details.
nthread	Integer. Number of threads to use for parallel computation during model training in XGBoost. Default is 1.
eps	Small value to avoid log(0) in LogLoss calculations. Default is 1e-15.
MC_sample	Not used by the model. The share of data per Monte Carlo sample is applied in test.gen before the wrapper is called. (xgboost's own row subsampling can be set with subsample in ...)
...	Additional arguments passed to xgb.train as parameters (e.g. eta, max_depth, objective).

Details

The type of task is decided from the metric and the response: "RMSE" gives regression, "Kappa" and "LogLoss" give classification (binary or multiclass). With a custom `metricfunc`, a numeric response gives regression and a factor, character or logical response gives classification. Classes are encoded as 0..K-1 internally, so any class labels work.

A custom `metricfunc` is called as `metricfunc(actual, predictions, ...)` (with the ... arguments only if the function accepts them), where `actual` has the type of the response (numeric for regression, factor for classification) and `predictions` is a numeric vector for regression, the probability of the second class level for binary classification, and an $n \times K$ probability matrix with the class levels as column names for multiclass classification.

Value

Performance metric

Index

* datasets

- ExponentialNoise, [15](#)
- HardCase, [17](#)
- NonLinearCategorization, [19](#)
- NonLinNormal, [20](#)
- NonLinNormalZs_d0, [21](#)
- NonLinNormalZs_d05, [22](#)
- NormalData, [23](#)
- PoissonNoise, [28](#)
- PolyData, [28](#)
- UniformNoise_large, [34](#)

add_interaction_terms, [2](#)

add_poly_terms, [3](#)

build_formula, [4](#)

CCI (CCI.test), [10](#)

CCI.direction, [5](#)

CCI.pretuner, [7](#), [13](#), [14](#)

CCI.test, [7](#), [10](#), [10](#), [17](#)

check_formula, [14](#)

clean_formula, [15](#)

ExponentialNoise, [15](#)

get_pvalues, [16](#)

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

HardCase, [17](#)

is_categorical_Z_any, [18](#)

make_strata_from_categorical_Z, [19](#)

NonLinearCategorization, [19](#)

NonLinNormal, [20](#)

NonLinNormalZs_d0, [21](#)

NonLinNormalZs_d05, [22](#)

NormalData, [23](#)

perm.test, [10](#), [14](#), [23](#), [27](#), [30](#), [31](#)

permute_within_strata, [26](#)

plot.CCI, [14](#), [25](#), [26](#), [27](#), [30](#), [31](#)

PoissonNoise, [28](#)

PolyData, [28](#)

print.CCI, [25](#), [27](#), [29](#), [31](#)

print.summary.CCI, [14](#)

print.summary.CCI (print.CCI), [29](#)

QQplot, [14](#), [25](#), [30](#), [30](#)

reports (print.CCI), [29](#)

summary.CCI, [25](#), [27](#), [31](#)

summary.CCI (print.CCI), [29](#)

test.gen, [25](#), [31](#), [31](#), [38](#)

tuner (CCI.pretuner), [7](#)

unclean_formula, [34](#)

UniformNoise_large, [34](#)

wrapper_knn, [7](#), [35](#)

wrapper_ranger, [7](#), [36](#)

wrapper_svm, [7](#), [37](#)

wrapper_xgboost, [7](#), [38](#)