

Package ‘liver’

April 20, 2021

Title “Eating the Liver of Data Science”

Version 1.7

Description Provides a collection of helper functions that make various techniques from data science more user-friendly for non-experts. In this way, our aim is to allow non-experts to become familiar with the techniques with only a minimal level of coding knowledge. Indeed, following an ancient Persian idiom, we refer to this as “eating the liver of data science” which could be interpreted as “getting intimately close with data science”. Examples of procedures we include are: data partitioning for out-of-sample testing, computing Mean Squared Error (MSE) for quantifying prediction accuracy, and data transformation (z-score and min-max). Besides such helper functions, the package also includes several interesting datasets that are useful for multivariate analysis.

URL <https://www.uva.nl/profile/a.mohammadi>

Depends R (>= 3.5.0)

Imports class, ggplot2

License GPL (>= 2)

Repository CRAN

Author Reza Mohammadi [aut, cre] (<<https://orcid.org/0000-0001-9538-0648>>),
Kevin Burke [aut]

Maintainer Reza Mohammadi <a.mohammadi@uva.nl>

NeedsCompilation no

Date/Publication 2021-04-20 15:20:02 UTC

R topics documented:

liver-package	2
accuracy	3
adult	4
advertising	5
bank	6
cereal	8
churn	9
churnTel	10

conf.mat	12
conf.mat.plot	13
fertilizer	14
find.na	14
house	15
housePrice	16
insurance	16
kNN	17
kNN.plot	19
mae	20
marketing	21
minmax	22
mse	23
partition	24
risk	24
skewness	25
transform	26
zscore	27
Index	28

liver-package

liver: "Eating the Liver of Data Science"

Description

The **liver** package provides a collection of helper functions that make various techniques from data science more user-friendly for non-experts. In this way, our aim is to allow non-experts to become familiar with the techniques with only a minimal level of coding knowledge. Indeed, following an ancient Persian idiom, we refer to this as "eating the liver of data science" which could be interpreted as "getting intimately close with data science". Examples of procedures we include are: data partitioning for out-of-sample testing, computing Mean Squared Error (MSE) for quantifying prediction accuracy, and data transformation (z-score and min-max). Besides such helper functions, the package also includes several interesting datasets that are useful for multivariate analysis.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl>
Amsterdam Business School
University of Amsterdam

Kevin Burke <kevin.burke@ul.ie>
Departement of Statistics
University of Limerick

Maintainer: Reza Mohammadi <a.mohammadi@uva.nl>

accuracy	<i>Average classification accuracy</i>
----------	--

Description

Computes average classification accuracy.

Usage

```
accuracy( pred, actual, cutoff = NULL, reference = NULL )
```

Arguments

pred	a numerical vector of estimated values.
actual	a numerical vector of actual values.
cutoff	cutoff value for the case that pred is vector of probabilities.
reference	a factor of classes to be used as the true results. .

Value

The computed average classification accuracy (numeric value).

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[conf.mat](#), [mse](#), [mae](#)

Examples

```
pred = c( "no", "yes", "yes", "no", "no", "yes", "no", "no" )
actual = c( "yes", "no", "yes", "no", "no", "no", "yes", "yes" )

accuracy( pred, actual )
```

adult

adult data set

Description

The adult dataset was collected from the US Census Bureau and the primary task is to predict whether a given adult makes more than \$50K a year based attributes such as education, hours of work per week, etc. The target feature is *income*, a factor with levels " $\leq 50K$ " and " $> 50K$ ", and the remaining 14 variables are predictors.

Usage

```
data( adult )
```

Format

The `adult` dataset, as a data frame, contains 48598 rows and 15 columns (variables/features). The 15 variables are:

- `age`: age in years.
- `workclass`: a factor with 6 levels.
- `demogweight`: the demographics to describe a person.
- `education`: a factor with 16 levels.
- `education.num`: number of years of education.
- `marital.status`: a factor with 5 levels.
- `occupation`: a factor with 15 levels.
- `relationship`: a factor with 6 levels.
- `race`: a factor with 5 levels.
- `gender`: a factor with levels "Female", "Male".
- `capital.gain`: capital gains.
- `capital.loss`: capital losses.
- `hours.per.week`: number of hours of work per week.
- `native.country`: a factor with 42 levels.
- `income`: yearly income as a factor with levels " $\leq 50K$ " and " $> 50K$ ".

Details

This dataset can be downloaded from the UCI machine learning repository:

<http://www.cs.toronto.edu/~delve/data/adult/desc.html>

A detailed description of the dataset can be found in the UCI documentation at:

<http://www.cs.toronto.edu/~delve/data/adult/adultDetail.html>

References

Kohavi, R. (1996). Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. *Kdd*.

See Also

[risk](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( adult )
```

```
str( adult )
```

advertising

advertising data set

Description

The dataset is from an anonymous organisation's social media ad campaign. The advertising dataset contains 11 features and 1143 records.

Usage

```
data( advertising )
```

Format

The advertising dataset, as a data frame, contains 1143 rows and 11 columns (variables/features). The 11 variables are:

- `ad.id`: an unique ID for each ad.
- `xyz.campaign.id`: an ID associated with each ad campaign of XYZ company.
- `fb.campaign.id`: an ID associated with how Facebook tracks each campaign.
- `age`: age of the person to whom the ad is shown.
- `gender`: gender of the person to whom the ad is shown.
- `interest`: a code specifying the category to which the person's interest belongs (interests are as mentioned in the person's Facebook public profile).
- `impressions`: the number of times the ad was shown.
- `clicks`: number of clicks on for that ad.
- `spend`: amount paid by company xyz to Facebook, to show that ad.
- `conversion`: total number of people who enquired about the product after seeing the ad.
- `approved`: total number of people who bought the product after seeing the ad.

Details

A detailed description of the dataset can be found:

<https://www.kaggle.com/loveall/clicks-conversion-tracking>

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [bank](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( advertising )
```

```
str( advertising )
```

bank

Bank marketing data set

Description

The data is related to direct marketing campaigns of a Portuguese banking institution. The marketing campaigns were based on phone calls. Often, more than one contact to the same client was required, in order to access if the product (bank term deposit) would be (or not) subscribed. The classification goal is to predict if the client will subscribe a term deposit (variable deposit).

Usage

```
data( bank )
```

Format

The bank dataset, as a data frame, contains 4521 rows (customers) and 17 columns (variables/features). The 17 variables are:

Bank client data:

- age: numeric.
- job: type of job; categorical: "admin.", "unknown", "unemployed", "management", "housemaid", "entrepreneur", "student", "blue-collar", "self-employed", "retired", "technician", "services".
- marital: marital status; categorical: "married", "divorced", "single"; note: "divorced" means divorced or widowed.
- education: categorical: "secondary", "primary", "tertiary", "unknown".
- default: has credit in default?; binary: "yes", "no".
- balance: average yearly balance, in euros; numeric.
- housing: has housing loan? binary: "yes", "no".
- loan: has personal loan? binary: "yes", "no".

Related with the last contact of the current campaign:

- `contact`: contact communication type; categorical: "unknown", "telephone", "cellular".
- `day`: last contact day of the month; numeric.
- `month`: last contact month of year; categorical: "jan", "feb", "mar", ..., "nov", "dec".
- `duration`: last contact duration, in seconds; numeric.

Other attributes:

- `campaign`: number of contacts performed during this campaign and for this client; numeric, includes last contact.
- `pdays`: number of days that passed by after the client was last contacted from a previous campaign; numeric, -1 means client was not previously contacted.
- `previous`: number of contacts performed before this campaign and for this client; numeric.
- `poutcome`: outcome of the previous marketing campaign; categorical: "success", "failure", "unknown", "other".

Target variable:

- `deposit`: Indicator of whether the client subscribed a term deposit; binary: "yes" or "no".

Details

This dataset can be downloaded from the UCI machine learning repository:

<http://archive.ics.uci.edu/ml/datasets/Bank+Marketing>

References

Moro, S., Laureano, R. and Cortez, P. (2011) Using Data Mining for Bank Direct Marketing: An Application of the CRISP-DM Methodology. In P. Novais et al. (Eds.), Proceedings of the European Simulation and Modelling Conference.

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( bank )
```

```
str( bank )
```

cereal

Cereal data set

Description

This dataset contains nutrition information for 77 breakfast cereals and includes 16 variables. The "rating" column is our target as a rating of the cereals (Possibly from Consumer Reports?).

Usage

```
data( cereal )
```

Format

The cereal dataset, as a data frame, contains 77 rows (breakfast cereals) and 16 columns (variables/features). The 16 variables are:

- name: Name of cereal.
- manuf: Manufacturer of cereal:
 - A: American Home Food Products;
 - G: General Mills;
 - K: Kelloggs;
 - N: Nabisco;
 - P: Post;
 - Q: Quaker Oats;
 - R: Ralston Purina;
- type: cold or hot.
- calories: calories per serving.
- protein: grams of protein.
- fat: grams of fat.
- sodium: milligrams of sodium.
- fiber: grams of dietary fiber.
- carbo: grams of complex carbohydrates.
- sugars: grams of sugars.
- potass: milligrams of potassium.
- vitamins: vitamins and minerals - 0, 25, or 100, indicating the typical percentage of FDA recommended.
- shelf: display shelf (1, 2, or 3, counting from the floor).
- weight: weight in ounces of one serving.
- cups: number of cups in one serving.
- rating: a rating of the cereals (Possibly from Consumer Reports?).

Details

These datasets have been gathered and cleaned up by Petra Isenberg, Pierre Dragicevic and Yvonne Jansen. The original source can be found: <https://perso.telecom-paristech.fr/eagan/class/igr204/datasets>

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [housePrice](#), [house](#)

Examples

```
data( cereal )
```

```
str( cereal )
```

churn

Churn data set

Description

This dataset comes from IBM Sample Data Sets. Customer *churn* occurs when customers stop doing business with a company, also known as customer attrition. The data set contains 5000 rows (customers) and 20 columns (features). The "Churn" column is our target which indicate whether customer churned (left the company) or not.

Usage

```
data( churn )
```

Format

The churn dataset, as a data frame, contains 5000 rows (customers) and 20 columns (variables/features). The 20 variables are:

- `state`: Categorical, for the 51 states and the District of Columbia.
- `area.code`: Categorical.
- `account.length`: count, how long account has been active.
- `voice.plan`: Categorical, yes or no, voice mail plan.
- `voice.messages`: Count, number of voice mail messages.
- `intl.plan`: Categorical, yes or no, international plan.
- `intl.mins`: Continuous, minutes customer used service to make international calls.
- `intl.calls`: Count, total number of international calls.
- `intl.charge`: Continuous, total international charge.
- `day.mins`: Continuous, minutes customer used service during the day.
- `day.calls`: Count, total number of calls during the day.

- `day.charge`: Continuous, total charge during the day.
- `eve.mins`: Continuous, minutes customer used service during the evening.
- `eve.calls`: Count, total number of calls during the evening.
- `eve.charge`: Continuous, total charge during the evening.
- `night.mins`: Continuous, minutes customer used service during the night.
- `night.calls`: Count, total number of calls during the night.
- `night.charge`: Continuous, total charge during the night.
- `customer.calls`: Count, number of calls to customer service.
- `churn`: Categorical, yes or no. Indicator of whether the customer has left the company (yes or no).

References

Larose, D. T. and Larose, C. D. (2014). Discovering knowledge in data: an introduction to data mining. *John Wiley & Sons*.

See Also

[adult](#), [risk](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( churn )  
  
str( churn )
```

churnTel

churnTel dataset

Description

Customer *churn* occurs when customers stop doing business with a company, also known as customer attrition. The data set contains 7043 rows (customers) and 21 columns (features). The "Churn" column is our target which indicate whether customer churned (left the company) or not.

Usage

```
data( churnTel )
```

Format

The churnTel dataset, as a data frame, contains 7043 rows (customers) and 21 columns (variables/features). The 21 variables are:

- `customer.ID`: Customer ID.
- `gender`: Whether the customer is a male or a female.
- `senior.citizen`: Whether the customer is a senior citizen or not (1, 0).
- `partner`: Whether the customer has a partner or not (yes, no).
- `dependent`: Whether the customer has dependents or not (yes, no).
- `tenure`: Number of months the customer has stayed with the company.
- `phone.service`: Whether the customer has a phone service or not (yes, no).
- `multiple.lines`: Whether the customer has multiple lines or not (yes, no, no phone service).
- `internet.service`: Customer's internet service provider (DSL, fiber optic, no).
- `online.security`: Whether the customer has online security or not (yes, no, no internet service).
- `online.backup`: Whether the customer has online backup or not (yes, no, no internet service).
- `device.protection`: Whether the customer has device protection or not (yes, no, no internet service).
- `tech.support`: Whether the customer has tech support or not (yes, no, no internet service).
- `streaming.TV`: Whether the customer has streaming TV or not (yes, no, no internet service).
- `streaming.movie`: Whether the customer has streaming movies or not (yes, no, no internet service).
- `contract`: The contract term of the customer (month to month, 1 year, 2 year).
- `paperless.bill`: Whether the customer has paperless billing or not (yes, no).
- `payment.method`: The customer's payment method (electronic check, mail check, bank transfer, credit card).
- `monthly.charge`: The amount charged to the customer monthly.
- `total.charges`: The total amount charged to the customer.
- `churn`: Whether the customer churned or not (yes or no).

Details

For more information related to the dataset see: <https://www.kaggle.com/blastchar/telco-customer-churn>

See Also

[adult](#), [risk](#), [churn](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( churnTel )
```

```
str( churnTel )
```

`conf.mat`*Confusion Matrix*

Description

Create a Confusion Matrix.

Usage

```
conf.mat( pred, actual, cutoff = NULL, reference = NULL,  
          proportion = FALSE, dnn = c( "Prediction", "Actual" ), ... )
```

Arguments

<code>pred</code>	a vector of estimated values.
<code>actual</code>	a vector of actual values.
<code>cutoff</code>	cutoff value for the case that <code>pred</code> is vector of probabilities.
<code>reference</code>	a factor of classes to be used as the true results.
<code>proportion</code>	Logical: FALSE (default) for a confusion matrix with number of cases. TRUE, for a confusion matrix with the proportion of cases.
<code>dnn</code>	the names to be given to the dimensions in the result (the dimnames names).
<code>...</code>	options to be passed to <code>table</code> .

Value

the results of `table` on `pred` and `actual`.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[conf.mat.plot](#), [accuracy](#)

Examples

```
pred = c( "no", "yes", "yes", "no", "no", "yes", "no", "no" )  
actual = c( "yes", "no", "yes", "no", "no", "no", "yes", "yes" )  
  
conf.mat( pred, actual )  
conf.mat( pred, actual, proportion = TRUE )
```

conf.mat.plot	<i>Plot Confusion Matrix</i>
---------------	------------------------------

Description

Plot a Confusion Matrix.

Usage

```
conf.mat.plot( pred, actual, cutoff = NULL, reference = NULL, conf.level = 0,
              margin = 1, color = c( "#ff83a8", "#83ff9b" ), ... )
```

Arguments

pred	a vector of estimated values.
actual	a vector of actual values.
cutoff	cutoff value for the case that pred is vector of probabilities.
reference	a factor of classes to be used as the true results.
conf.level	confidence level used for the confidence rings on the odds ratios. Must be a single nonnegative number less than 1; if set to 0 (the default), confidence rings are suppressed.
margin	a numeric vector with the margins to equate. Must be one of 1 (the default), 2, or c(1, 2), which corresponds to standardizing the row, column, or both margins in each 2 by 2 table. Only used if std equals "margins".
color	a vector of length 2 specifying the colors to use for the smaller and larger diagonals of each 2 by 2 table.
...	options to be passed to fourfoldplot.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[conf.mat](#)

Examples

```
pred = c( "no", "yes", "yes", "no", "no", "yes", "no", "no" )
actual = c( "yes", "no", "yes", "no", "no", "no", "yes", "yes" )

conf.mat.plot( pred, actual )
```

fertilizer

Fertilizer data set

Description

The fertilizer dataset contains 4 features and 96 records. Results from an experiment to compare yields of a crop obtained under three different fertilizers. The target feature is *yield*.

Usage

```
data( fertilizer )
```

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( fertilizer )  
  
str( fertilizer )
```

find.na

find.na

Description

Finding missing values.

Usage

```
find.na( x )
```

Arguments

`x` a numerical vector, matrix or `data.frame`.

Value

A numeric matrix with two columns.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

Examples

```
x = c( 2.3, NA, -1.4, 0, 3.45 )  
  
find.na( x )
```

house	<i>house data set</i>
-------	-----------------------

Description

The house dataset contains 6 features and 414 records. The target feature is *unit.price* and the remaining 5 variables are predictors.

Usage

```
data( house )
```

Format

The house dataset, as a data frame, contains 414 rows and 6 columns (variables/features). The 6 variables are:

- `house.age`: house age (numeric, in year).
- `distance.to.MRT`: distance to the nearest MRT station (numeric).
- `stores.number`: number of convenience stores (numeric).
- `latitude`: latitude (numeric).
- `longitude`: longitude (numeric).
- `unit.price`: house price of unit area (numeric).

Details

A detailed description of the dataset can be found:

<https://www.kaggle.com/quantbruce/real-estate-price-prediction>

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#)

Examples

```
data( house )  
  
str( house )
```

`housePrice`*housePrice dataset*

Description

This data set contains 1460 rows and 81 columns (features). The "SalePrice" column is the target.

Usage

```
data( housePrice )
```

Format

The housePrice dataset, as a data frame, contains 1460 rows and 81 columns (variables/features).

Details

For more information related to the dataset see:

<https://www.kaggle.com/c/house-prices-advanced-regression-techniques/data>

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [house](#)

Examples

```
data( housePrice )
```

```
str( housePrice )
```

`insurance`*insurance data set*

Description

The insurance dataset contains 7 features and 1338 records. The target feature is *charge* and the remaining 6 variables are predictors.

Usage

```
data( insurance )
```


Format

The insurance dataset, as a data frame, contains 1338 rows (customers) and 7 columns (variables/features). The 7 variables are:

- age: age of primary beneficiary.
- bmi: body mass index, providing an understanding of body, weights that are relatively high or low relative to height, objective index of body weight (kg / m ^ 2) using the ratio of height to weight, ideally 18.5 to 24.9.
- children: Number of children covered by health insurance / Number of dependents.
- smoker: Smoking as a factor with 2 levels, yes, no.
- gender: insurance contractor gender, female, male.
- region: the beneficiary's residential area in the US, northeast, southeast, southwest, northwest.
- charge: individual medical costs billed by health insurance.

Details

A detailed description of the dataset can be found: <https://www.kaggle.com/mirichoi0218/insurance>

References

Brett Lantz (2019). Machine Learning with R: Expert techniques for predictive modeling. *Packt Publishing Ltd.*

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( insurance )  
  
str( insurance )
```

kNN

k-Nearest Neighbour Classification

Description

kNN is used to perform k-nearest neighbour classification for test set using training set. For each row of the test set, the k nearest (based on Euclidean distance) training set vectors are found. Then, the classification is done by majority vote (ties broken at random). This function provides a formula interface to the [knn](#) function of R package `class`. In addition, it allows normalization of the given data using the [transform](#) function.

Usage

```
kNN( formula, train, test, k = 1, transform = FALSE, l = 0, prob = FALSE,
      use.all = TRUE, na.rm = FALSE )
```

Arguments

formula	a formula , with a response but no interaction terms. For the case of data frame, it is taken as the model frame (see model.frame).
train	data frame or matrix of train set cases.
test	data frame or matrix of test set cases.
k	number of neighbours considered.
transform	a character with options FALSE (default), "minmax", and "zscore". Option "minmax" means no transformation. This option allows the users to use normalized version of the train and test sets for the kNN algorithm.
l	minimum vote for definite decision, otherwise doubt. (More precisely, less than k-1 dissenting votes are allowed, even if k is increased by ties.)
prob	If this is true, the proportion of the votes for the winning class are returned as attribute prob.
use.all	controls handling of ties. If true, all distances equal to the kth largest are included. If false, a random selection of distances equal to the kth is chosen to use exactly k neighbours.
na.rm	a logical value indicating whether NA values in x should be stripped before the computation proceeds.

Value

Factor of classifications for the test set, in which the doubt will be returned as NA; basically, the return value is the same as in the [knn](#) function of R package `class`.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

References

- Ripley, B. D. (1996) *Pattern Recognition and Neural Networks*. Cambridge.
- Venables, W. N. and Ripley, B. D. (2002) *Modern Applied Statistics with S*. Fourth edition. Springer.

See Also

[knn](#), [transform](#)

Examples

```
data( risk )

train = risk[ 1:100, ]
test  = risk[ 101, ]

kNN( risk ~ income + age, train = train, test = test )
```

kNN.plot

*Visualizing the Optimal Number of k***Description**

Visualizing the Optimal Number of k for k-Nearest Neighbour Classification kNN based on accuracy or Mean Square Error (MSE).

Usage

```
kNN.plot( formula, train, test, k.max = 10, transform = FALSE, base = "error",
          set.seed = NULL, ... )
```

Arguments

formula	a formula , with a response but no interaction terms. For the case of data frame, it is taken as the model frame (see model.frame).
train	data frame or matrix of train set cases.
test	data frame or matrix of test set cases.
k.max	the maximum number of number of neighbours to consider, must be at least two.
transform	a character with options FALSE (default), "minmax", and "zscore". Option "minmax" means no transformation. This option allows the users to use normalized version of the train and test sets for the kNN algorithm.
base	base measurement: error (default), accuracy, or MSE for Mean Square Error.
set.seed	a single value, interpreted as an integer, or NULL.
...	options to be passed to kNN().

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

References

Ripley, B. D. (1996) *Pattern Recognition and Neural Networks*. Cambridge.
 Venables, W. N. and Ripley, B. D. (2002) *Modern Applied Statistics with S*. Fourth edition. Springer.

See Also

[kNN](#), [transform](#)

Examples

```
data( risk )

train = risk[ 1:150, ]
test  = risk[ 151:246, ]

kNN.plot( risk ~ income + age, train = train, test = test )
kNN.plot( risk ~ income + age, train = train, test = test, base = "accuracy" )
```

mae	<i>Mean Absolute Error (MAE)</i>
-----	----------------------------------

Description

Computes mean absolute error.

Usage

```
mae( pred, actual, weight = 1, na.rm = FALSE )
```

Arguments

pred	a numerical vector of estimated values.
actual	a numerical vector of actual values.
weight	a numerical vector of weights the same length as pred.
na.rm	a logical value indicating whether NA values in pred should be stripped before the computation proceeds.

Value

The computed mean squared error (numeric value).

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[mse](#)

Examples

```
pred = c( 2.3, -1.4, 0, 3.45 )
actual = c( 2.1, -0.9, 0, 2.99 )

mae( pred, actual )
```

marketing

marketing data set

Description

The marketing dataset contains 8 features and 40 records as 40 days that report how much we spent, how many clicks, impressions and transactions we got, whether or not a display campaign was running, as well as our revenue, click-through-rate and conversion rate. The target feature is *revenue* and the remaining 7 variables are predictors.

Usage

```
data( marketing )
```

Format

The marketing dataset, as a data frame, contains 40 rows and 8 columns (variables/features). The 8 variables are:

- `spend`: daily spend of money on PPC (pay-per-click).
- `clicks`: number of clicks on for that ad.
- `impressions`: amount of impressions per day.
- `display`: whether or not a display campaign was running.
- `transactions`: number of transactions per day.
- `click.rate`: click-through-rate.
- `conversion.rate`: conversion rate.
- `revenue`: daily revenue.

Details

A detailed description of the dataset can be found:

<https://github.com/chrisBow/marketing-regression-part-one>

See Also

[adult](#), [risk](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( marketing )

str( marketing )
```

minmax	<i>Min-Max normalization</i>
--------	------------------------------

Description

Performs Min-Max normalization of numerical variables.

Usage

```
minmax( x, columns = NULL, na.rm = FALSE )
```

Arguments

x	a numerical vector, matrix or data.frame.
columns	which columns are going to transfer for the cases that x is a matrix or a data.frame. Defaults to all columns.
na.rm	a logical value indicating whether NA values in x should be stripped before the computation proceeds.

Value

transformed version of x.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[transform](#), [zscore](#)

Examples

```
x = c( 2.3, -1.4, 0, 3.45 )  
  
minmax( x )
```

mse	<i>Mean Squared Error (MSE)</i>
-----	---------------------------------

Description

Computes mean squared error.

Usage

```
mse( pred, actual, weight = 1, na.rm = FALSE )
```

Arguments

pred	a numerical vector of estimated values.
actual	a numerical vector of actual values.
weight	a numerical vector of weights the same length as pred.
na.rm	a logical value indicating whether NA values in pred should be stripped before the computation proceeds.

Value

The computed mean squared error (numeric value).

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[mae](#)

Examples

```
pred = c( 2.3, -1.4, 0, 3.45 )  
actual = c( 2.1, -0.9, 0, 2.99 )  
  
mse( pred, actual )
```

partition	<i>Partition the data</i>
-----------	---------------------------

Description

Randomly partitions the data (primarily intended to split into "training" and "test" sets) according to the supplied probabilities.

Usage

```
partition( data, prob = c( 0.7, 0.3 ), set.seed = NULL )
```

Arguments

data	an $(n \times p)$ matrix or a data.frame.
prob	a numerical vector in [0, 1].
set.seed	a single value, interpreted as an integer, or NULL.

Value

a list which includes the data partitions.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

Examples

```
data( iris )  
  
partition( data = iris, prob = c( 0.7, 0.3 ) )
```

risk	<i>Risk data set</i>
------	----------------------

Description

The *risk* dataset containing 6 features and 246 records. The target feature is *risk*, a factor with levels "good risk" and "bad risk" along with 5 predictors.

Usage

```
data( risk )
```


Format

The risk dataset, as a data frame, contains 246 rows (customers) and 6 columns (variables/features). The 6 variables are:

- age: age in years.
- marital: A factor with levels "single", "married", and "other".
- income: yearly income.
- mortgage: A factor with levels "yes" and "no".
- nr_loans: Number of loans that constomers have.
- risk: A factor with levels "good risk" and "bad risk".

See Also

[adult](#), [churn](#), [churnTel](#), [bank](#), [advertising](#), [marketing](#), [insurance](#), [cereal](#), [housePrice](#), [house](#)

Examples

```
data( risk )  
  
str( risk )
```

skewness

Skewness

Description

Computes the skewness for each field.

Usage

```
skewness( x, na.rm = FALSE )
```

Arguments

x a numerical vector, matrix or data.frame.
na.rm a logical value indicating whether NA values in x should be stripped before the computation proceeds.

Value

A numeric vector of skewness values.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

Examples

```
x = c( 2.3, -1.4, 0, 3.45 )  
  
skewness( x )
```

transform

Z-score normalization

Description

Performs variable transformation such as Z-score and min-max normalization.

Usage

```
transform( x, method = c( "minmax", "zscore" ), columns = NULL, na.rm = FALSE )
```

Arguments

x	a numerical vector, a matrix or a data.frame.
method	a method to transfer x.
columns	which columns are going to transfer for the cases that x is a matrix or a data.frame. Defaults to all columns.
na.rm	a logical value indicating whether NA values in x should be stripped before the computation proceeds.

Value

transformed version of x.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[zscore](#), [minmax](#)

Examples

```
x = c( 2.3, -1.4, 0, 3.45 )  
  
transform( x, method = "minmax" )  
  
transform( x, method = "zscore" )
```

zscore	<i>Z-score normalization</i>
--------	------------------------------

Description

Performs Z-score normalization to transform numerical variables.

Usage

```
zscore( x, columns = NULL, na.rm = FALSE )
```

Arguments

x	a numerical vector, matrix or data.frame.
columns	which columns are going to transfer for the cases that x is a matrix or a data.frame. Defaults to all columns.
na.rm	a logical value indicating whether NA values in x should be stripped before the computation proceeds.

Value

transformed version of x.

Author(s)

Reza Mohammadi <a.mohammadi@uva.nl> and Kevin Burke <kevin.burke@ul.ie>

See Also

[transform](#), [minmax](#)

Examples

```
x = c( 2.3, -1.4, 0, 3.45 )  
  
zscore( x )
```

Index

* data preprocessing

find.na, 14
minmax, 22
partition, 24
skewness, 25
transform, 26
zscore, 27

* datasets

adult, 4
advertising, 5
bank, 6
cereal, 8
churn, 9
churnTel, 10
fertilizer, 14
house, 15
housePrice, 16
insurance, 16
marketing, 21
risk, 24

* models

kNN, 17
kNN.plot, 19

* package

liver-package, 2

* parameter learning

accuracy, 3
conf.mat, 12
conf.mat.plot, 13
mae, 20
mse, 23
skewness, 25

accuracy, 3, 12

adult, 4, 6, 7, 9–11, 14–17, 21, 25

advertising, 5, 5, 7, 9–11, 14–17, 21, 25

bank, 5, 6, 6, 9–11, 14–17, 21, 25

cereal, 5–7, 8, 10, 11, 14–17, 21, 25

churn, 5–7, 9, 9, 11, 14–17, 21, 25

churnTel, 5–7, 9, 10, 10, 14–17, 21, 25

conf.mat, 3, 12, 13

conf.mat.plot, 12, 13

fertilizer, 14

find.na, 14

formula, 18, 19

house, 5–7, 9–11, 14, 15, 16, 17, 21, 25

housePrice, 5–7, 9–11, 14, 15, 16, 17, 21, 25

insurance, 5–7, 9–11, 14–16, 16, 21, 25

kNN, 17, 20

knn, 17, 18

kNN.plot, 19

liver-package, 2

mae, 3, 20, 23

marketing, 5–7, 9–11, 14–17, 21, 25

minmax, 22, 26, 27

model.frame, 18, 19

mse, 3, 20, 23

partition, 24

risk, 5–7, 9–11, 14–17, 21, 24

skewness, 25

transform, 17, 18, 20, 22, 26, 27

zscore, 22, 26, 27