

Package ‘irrCAC’

August 25, 2019

Author Kilem L. Gwet, Ph.D.

Version 1.0

Date 2018-02-14

Maintainer Kilem L. Gwet <gwet@agreestat.com>

Title Computing the extent of agreement among raters with chance-corrected agreement coefficient (CAC)

Description

This package contains a series of R functions for calculating various chance-corrected agreement coefficients (CAC) among 2 or more raters. Among the CAC coefficients covered are Cohen's kappa, Conger's kappa, Fleiss' kappa, Brennan-Prediger coefficient, Gwet's AC1/AC2 coefficients, and Krippendorff's alpha. Multiple sets of weights are proposed for computing weighted analyses.

License GPL (>=2)

Encoding UTF-8

LazyData true

RoxygenNote 6.1.1

Suggests knitr,
rmarkdown

VignetteBuilder knitr

R topics documented:

altman	2
altman.bf	3
bipolar.weights	3
bp.coeff.dist	4
bp.coeff.raw	5
bp2.table	6
cac.ben.gerry	6
cac.dist.g1g2	7
cac.dist4cat	8
cac.raw.g1g2	8
cac.raw.gender	9
cac.raw4raters	9
cac.raw5obser	10

circular.weights	11
conger.kappa.raw	11
cont3x3abstractors	12
cont4x4diagnosis	13
distrib.6raters	13
fleiss	14
fleiss.bf	15
fleiss.kappa.dist	15
fleiss.kappa.raw	16
gwet.ac1.dist	17
gwet.ac1.raw	18
gwet.ac1.table	19
identity.weights	20
kappa2.table	20
krippen.alpha.dist	21
krippen.alpha.raw	22
krippen2.table	23
landis.koch	24
landis.koch.bf	24
linear.weights	25
ordinal.weights	25
pa.coeff.dist	26
pa.coeff.raw	27
pa2.table	28
quadratic.weights	28
radical.weights	29
ratio.weights	29
scott2.table	30
trim	30

Index **31**

altman	<i>Dataset describing the Altman's Benchmarking Scale</i>
--------	---

Description

This dataset contains information describing the Altman scale for benchmarking chance-corrected agreement coefficients such as Gwet AC1/AC2, Kappa and many others.

Usage

altman

Format

Each row of this dataset describes an interval and the interpretation of the magnitude it represents.

lb.AL The interval lower bound

ub.AL The interval upper bound

interp.AL The interpretation

Source

Altman, D.G. (1991). *Practical Statistics for Medical Research*. Chapman and Hall.

altman.bf

Computing Altman's Benchmark Scale Membership Probabilities

Description

Computing Altman's Benchmark Scale Membership Probabilities

Usage

```
altman.bf(coeff, se, BenchDF = altman)
```

Arguments

coeff	A mandatory parameter representing the estimated value of an agreement coefficient.
se	A mandatory parameter representing the agreement coefficient standard error.
BenchDF	An optional parameter that is a 3-column data frame containing the Altman's benchmark scale information. The 3 columns are the interval lower bound, upper bound, and their interpretation. The default value is a small file contained in the package and named <i>altman.RData</i> , which describes the official Altman's scale intervals and their interpretation.

Value

A one-column matrix containing the membership probabilities (c.f. http://agreestat.com/research_papers/inter-rater%20reliability%20study%20design1.pdf)

bipolar.weights

Function for computing the Bipolar Weights

Description

Function for computing the Bipolar Weights

Usage

```
bipolar.weights(categ)
```

Arguments

categ	A mandatory parameter representing the vector of all possible ratings.
-------	--

Value

A square matrix of quadratic weights to be used for calculating the weighted coefficients.

bp.coeff.dist	<i>Brennan-Prediger's agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.</i>
---------------	---

Description

Brennan-Prediger's agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.

Usage

```
bp.coeff.dist(ratings, weights = "unweighted", categ = NULL,
             conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxq matrix / data frame containing the distribution of raters by subject and category. Each cell (i,k) contains the number of raters who classified subject i into category k .
weights	is an optional parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter <code>categ</code> . Otherwise, only the categories reported will be used.
categ	An optional parameter representing all categories available to raters during the experiment. This parameter may be useful if some categories were not used by any rater in spite of being available to the raters.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A vector containing the following information: pa(the percent agreement),pe(the percent chance agreement),coeff(Brennan-Prediger coefficient), stderr(the standard error of Brennan-Prediger coefficient),conf.int(the p-value of Brennan-Prediger coefficient), p.value(the p-value of Brennan-Prediger coefficient),coeff.name ("Brennan-Prediger").

Source

Brennan, R.L., and Prediger, D. J. (1981). "Coefficient Kappa: some uses, misuses, and alternatives," *Educational and Psychological Measurement*, 41, 687-699.

bp.coeff.raw	<i>Brennan & Prediger's agreement coefficient for an arbitrary number of raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.</i>
--------------	---

Description

Brennan & Prediger's agreement coefficient for an arbitrary number of raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.

Usage

```
bp.coeff.raw(ratings, weights = "unweighted", categ.labels = NULL,
             conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxr matrix / data frame of ratings where each column represents one rater and each row one subject.
weights	is a mandatory parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter categ.labels. Otherwise, the program may not work.
categ.labels	An optional vector parameter containing the list of all possible ratings. It may be useful in case some of the possible ratings are not used by any rater, they will still be used when calculating agreement coefficients. The default value is NULL. In this case, only categories reported by the raters are used in the calculations.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A data list containing 3 objects: (1) a one-row data frame containing various statistics including the requested agreement coefficient, (2) the weight matrix used in the calculations if any, and (3) A vector of categories used in the analysis. These could be categories reported by the raters, or those available to the raters whether they used them or not. The output data frame contains the following variables: "coeff.name" (coefficient name), "pa" (the percent agreement), "pe" (the percent chance agreement), "coeff.val" (Brennan-Prediger coefficient estimate), "coeff.se" (standard error), "conf.int" (the confidence interval), "p.value" (Brennan-Prediger coefficient's p-value), "w.name" (the weights' identification).

References

Brennan, R.L., & Prediger, D. J. (1981). "Coefficient Kappa: some uses, misuses, and alternatives." *Educational and Psychological Measurement*, 41, 687-699.

bp2.table	<i>Brenann-Prediger coefficient for 2 raters</i>
-----------	--

Description

Brenann-Prediger coefficient for 2 raters

Usage

```
bp2.table(ratings, weights = identity.weights(1:ncol(ratings)),
  conflev = 0.95, N = Inf)
```

Arguments

ratings	A square table of ratings (assume no missing ratings).
weights	An optional matrix that contains the weights used in the weighted analysis. By default, this parameter contains the identity weight matrix, which leads to the unweighted analysis.
conflev	An optional parameter that specifies the confidence level used for constructing confidence intervals. By default the function assumes the standard value of 95%.
N	An optional parameter representing the finite population size if any. It is used to perform the finite population correction to the standard error. It's default value is infinity.

Value

A data frame containing the following 5 variables: coeff.name coeff.val coeff.se coeff.ci coeff.pval.

cac.ben.gerry	<i>Ratings of 12 units from 2 raters named Ben and Gerry</i>
---------------	--

Description

This dataset contains ratings that 2 raters named Ben and Gerry assigned to 12 units distributed in 2 groups "G1" and "G2".

Usage

```
cac.ben.gerry[,c(3,4)] or cac.ben.gerry[,c("Ben", "Gerry")]
```

Format

Each row of this dataset describes an interval and the interpretation of the magnitude it represents.

Group Group Name

Units Unit number

Ben Ben's Ratings

Gerry Gerry's Ratings

The first 2 columns "Group" and "Units" play a descriptive role here and are not used by any function included in this package. One will typically use `cac.ben.gerry[,c(3,4)]` or `cac.ben.gerry[,c("Ben", "Gerry")]` as input dataset.

`cac.dist.g1g2`

Distribution of 4 raters by subject and by category, for 14 Subjects that belong to 2 groups "G1" and "G2"

Description

This dataset contains rating data in the form of a subject-level distribution of 4 raters by category the subject was classified into. A total of 4 raters had to classify 14 subjects into one of 5 categories "a", "b", "c", "d", and "e". This dataset is different version of the more detailed `cac.raw.g1g2` dataset. While `cac.raw.g1g2` tells you about the exact category into which each rater classified all subjects, `cac.dist.g1g2` on the other hand, can only tell you how many raters classified a given subject into a particular category.

Usage

```
cac.raw.g1g2[,3:6]
```

Format

This dataset contains ratings obtained from an experiment where 4 raters classified 14 subjects into 5 possible categories labeled as a, b, c, d, and e. None of the 4 raters scored all 14 units. Therefore, some missing ratings appear in each of the columns associated with the 4 raters.

Note that only the the 4 last columns are to be used with the functions included in this package. The first 2 columns only play a descriptive role and are not used in any calculation.

Group This variable represents the group name.

Units This variable represents the unit number.

a Number of raters who classified the subject represented by the row into category "a"

b Number of raters who classified the subject represented by the row into category "b"

c Number of raters who classified the subject represented by the row into category "c"

d Number of raters who classified the subject represented by the row into category "d"

e Number of raters who classified the subject represented by the row into category "e"

cac.dist4cat	<i>Distribution of 4 raters by Category and Subject - Subjects allocated in 2 groups A and B.</i>
--------------	---

Description

This dataset summarizes the ratings assigned by 4 raters who classified 15 subjects into one of 3 categories named "a", "b", and "c".

Usage

cac.dist4cat

Format

This dataset has 15 rows (for the 15 subjects) and 4 columns. Only the last 3 columns representing the categories into which subjects are classified are used in the calculations - unless the sub-group analysis is required.

Group This variable represents the subject number.

a category a

b Category b

c Category c

cac.raw.g1g2	<i>Dataset of raw ratings from 4 Raters on 14 Subjects that belong to 2 groups named "G1" and "G2"</i>
--------------	--

Description

This dataset contains data from a reliability experiment where 4 raters identified as Rater1, Rater2, Rater3 and Rater4 scored 14 units on a 5-point alphabetical scale based on the values a, b, c, d and e. These 14 units are allocated to 2 groups named G1 and G2.

Usage

cac.raw.g1g2

Format

This dataset contains ratings obtained from an experiment where 4 raters classified 14 subjects into 5 possible categories labeled as a, b, c, d, and e. None of the 4 raters scored all 14 units. Therefore, some missing ratings appear in each of the columns associated with the 4 raters.

Note that only the the 4 last columns are to be used with the functions included in this package. The first 2 columns only play a descriptive role and are not used in any calculation.

Group This variable represents the unit number.

Units This variable represents the unit number.

Rater1 All ratings from rater 1

Rater2 All ratings from rater 2

Rater3 All ratings from rater 3

Rater4 All ratings from rater 4

cac.raw.gender	<i>Rating Data from 4 Raters and 15 human Subjects, 9 of whom are female and 6 males.</i>
----------------	---

Description

This dataset contains data from a reliability experiment where 4 raters scored 15 units on a 3-point alphabetic scale based on the values a, b, and c.

Usage

cac.raw.gender

Format

This dataset contains ratings obtained from an experiment where 4 raters classified 15 subjects into 3 possible categories labeled as a, b, and c.

Note that only the the 4 last columns are to be used with the functions included in this package. The first column only plays a descriptive role and is not to be used in any calculation.

Group This variable represents the unit number.

RaterA All ratings from rater 1

RaterB All ratings from rater 2

RaterC All ratings from rater 3

RaterD All ratings from rater 4

cac.raw4raters	<i>Rating Data from 4 Raters and 12 Subjects.</i>
----------------	---

Description

This dataset contains data from a reliability experiment where 5 observers scored 15 units on a 4-point numeric scale based on the values 0, 1, 2 and 3.

Usage

cac.raw4raters

Format

This dataset contains ratings obtained from an experiment where 4 raters classified 12 subjects into 5 possible categories labeled as 1, 2, 3, 4, and 5. None of the 4 raters scored all 12 units. Therefore, some missing ratings in the form of "NA" appear in each of the columns associated with the 4 raters.

Note that only the the 4 last columns are to be used with the functions included in this package. The first column only plays a descriptive role and is not used in any calculation.

Units This variable represents the unit number.

Rater1 All ratings from rater 1

Rater2 All ratings from rater 2

Rater3 All ratings from rater 3

Rater4 All ratings from rater 4

Source

Gwet, K.L. (2014) *Handbook of Inter-Rater Reliability*, 4th Edition, page #120. Advanced Analytics, LLC.

cac.raw5obser

Scores assigned by 5 observers to 20 experimental units.

Description

This dataset contains data from a reliability experiment where 5 observers scored 15 units on a 4-point numeric scale based on the values 0, 1, 2 and 3.

Usage

cac.raw5obser

Format

This dataset has 15 rows (for the 15 subjects) and 6 columns. Only the last 5 columns associated with the 5 observers are used in the calculations. Of the 5 observers, only observer 3 scored all 15 units. Therefore, some missing ratings in the form of "NA" appear in the columns associated with the remaining 4 observers.

Unit This variable represents the unit number.

Observer1 All ratings from Observer 1

Observer2 All ratings from Observer 2

Observer3 All ratings from Observer 3

Observer4 All ratings from Observer 4

Observer5 All ratings from Observer 5

Source

Gwet, K.L. (2014) *Handbook of Inter-Rater Reliability*, 4th Edition. Advanced Analytics, LLC. A larger version of this table can be found on page #125

circular.weights	<i>Function for computing the Circular Weights</i>
------------------	--

Description

Function for computing the Circular Weights

Usage

```
circular.weights(categ)
```

Arguments

categ	A mandatory parameter representing the vector of all possible ratings.
-------	--

Value

A square matrix of quadratic weights to be used for calculating the weighted coefficients.

conger.kappa.raw	<i>Conger's generalized kappa coefficient for an arbitrary number of raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.</i>
------------------	---

Description

Conger's generalized kappa coefficient for an arbitrary number of raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.

Usage

```
conger.kappa.raw(ratings, weights = "unweighted", categ.labels = NULL,
  conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxr matrix / data frame of ratings where each column represents one rater and each row one subject.
weights	is a mandatory parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter categ.labels. Otherwise, the program may not work.
categ.labels	An optional vector parameter containing the list of all possible ratings. It may be useful in case some of the possible ratings are not used by any rater, they will still be used when calculating agreement coefficients. The default value is NULL. In this case, only categories reported by the raters are used in the calculations.

conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A data list containing 3 objects: (1) a one-row data frame containing various statistics including the requested agreement coefficient, (2) the weight matrix used in the calculations if any, and (3) A vector of categories used in the analysis. These could be categories reported by the raters, or those available to the raters whether they used them or not. The output data frame contains the following variables: "coeff.name" (coefficient name), "pa" (the percent agreement), "pe" (the percent chance agreement), "coeff.val" (Conger's Kappa estimate), "coeff.se" (standard error), "conf.int" (Conger Kappa's confidence interval), "p.value" (agreement coefficient's p-value), "w.name" (the weights' identification).

References

Conger, A. J. (1980), "Integration and Generalization of Kappas for Multiple Raters," *Psychological Bulletin*, 88, 322-328.

cont3x3abstractors	<i>Distribution of 100 pregnant women by pregnancy type and by abstractor.</i>
--------------------	--

Description

This dataset contains pregnancy type data collected from 100 women who entered an Emergency Room with a positive pregnancy test and a second condition, which is either abdominal pain or vaginal bleeding. After reviewing their medical records, 2 reviewers (also referred to as abstractors) classified them into one of the following three pregnancy categories: Ectopic Pregnancy (Ectopic), Abnormal Intrauterine pregnancy (AIU) and Normal Intrauterine Pregnancy (NIU).

Usage

cont3x3abstractors

Format

Each row of this dataset describes an interval and the interpretation of the magnitude it represents.

Type Pregnancy Type. This variable is shown here for information only and is never used by any function in the irrCAC package.

Ectopic Ectopic Pregnancy

AIU Abnormal Intrauterine Pregnancy

NIU Normal Intrauterine Pregnancy

Source

Gwet, K.L. (2014). *Handbook of Inter-Rater Reliability*, 4th Edition. Advanced Analytics, LLC.

cont4x4diagnosis	<i>Distribution of 223 Psychiatric Patients by Type of of Psychiatric Disorder and Diagnosis Method.</i>
------------------	--

Description

This dataset shows the distribution of 223 psychiatric patients by diagnosis category and by the method used to obtain the diagnosis. The first method named "Clinical Diagnosis" (also known as "Facility Diagnosis") is used in a service facility (e.g. public hospital, or a community unit) and does not rely on a rigorous application of research criteria. The second method known as "Research Diagnosis" is based on a strict application of research criteria. Column 1 contains the diagnosis categories into which patients are classified with Method 1. The first row on the other hand, shows categories into which patients are classified with Method 2.

Usage

cont4x4diagnosis

Format

This dataset contains a 4x4 squared table. The first column is never used in the calculations and only contains row names. Only the last 4 columns are used for computing agreement coefficients.

Diagnosis Pregnancy Type. This variable is shown here for information only and is never used by any function in the irrCAC package.

Schizophrenia Ectopic Pregnancy

Bipolar.Disorder Abnormal Intrauterine Pregnancy

Depression Normal Intrauterine Pregnancy

Other Normal Intrauterine Pregnancy

Source

Gwet, K.L. (2014). *Handbook of Inter-Rater Reliability*, 4th Edition. Advanced Analytics, LLC.

distrib.6raters	<i>Distribution of 6 psychiatrists by Subject/patient and diagnosis Category.</i>
-----------------	---

Description

This dataset summarizes the ratings assigned by 6 psychiatrists classifying 15 patients into one of five categories named "Depression", "Personal Disorder", "Schizophrenia", "Neurosis" and "Other".

Usage

distrib.6raters

Format

This dataset has 15 rows (for the 15 subjects) and 7 columns. Only the last 6 columns representing the categories into which subjects are classified are used in the calculations.

Subject This variable represents the subject number.

Personality.Disorder Personality disorder category

Schizophrenia Schizophrenia Category

Neurosis Neurosis category

Other "Other" category

Source

Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters, *Psychological Bulletin*, 76, 378-382.

 fleiss

Dataset describing Fleiss' Benchmarking Scale

Description

This dataset contains information describing Fleiss' scale for benchmarking chance-corrected agreement coefficients such as Gwet AC1/AC2, Kappa and many others.

Usage

fleiss

Format

Each row of this dataset describes an interval and the interpretation of the magnitude it represents.

lb.FL The interval lower bound

ub.FL The interval upper bound

interp.FL The interpretation

Source

Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. John Wiley & Sons.

fleiss.bf

Computing Fleiss Benchmark Scale Membership Probabilities

Description

Computing Fleiss Benchmark Scale Membership Probabilities

Usage

```
fleiss.bf(coeff, se, BenchDF = fleiss)
```

Arguments

coeff	A mandatory parameter representing the estimated value of an agreement coefficient.
se	A mandatory parameter representing the agreement coefficient standard error.
BenchDF	An optional parameter that is a 3-column data frame containing the Fleiss' benchmark scale information. The 3 columns are the interval lower bound, upper bound, and their interpretation. The default value is a small file contained in the package and named <i>fleiss.RData</i> , which describes the fleiss' scale intervals and their interpretation.

Value

A one-column matrix containing the membership probabilities (c.f. http://agreestat.com/research_papers/inter-rater%20reliability%20study%20design1.pdf)

fleiss.kappa.dist

Fleiss' agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.

Description

Fleiss' agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.

Usage

```
fleiss.kappa.dist(ratings, weights = "unweighted", categ = NULL,
  conflev = 0.95, N = Inf)
```

Arguments

ratings	An $n \times q$ matrix / data frame containing the distribution of raters by subject and category. Each cell (i,k) contains the number of raters who classified subject i into category k .
---------	---

weights	is an optional parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix $q \times q$ where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter <code>categ</code> . Otherwise, only the categories reported will be used.
categ	An optional parameter representing all categories available to raters during the experiment. This parameter may be useful if some categories were not used by any rater in spite of being available to the raters.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A vector containing the following information: `pa`(the percent agreement), `pe`(the percent chance agreement), `coeff`(Fleiss' agreement coefficient), `stderr`(the standard error of Fleiss' coefficient), `conf.int`(the confidence interval of Fleiss Kappa coefficient), `p.value`(the p-value of Fleiss' coefficient), `coeff.name` ("Fleiss").

Source

Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. John Wiley & Sons.

fleiss.kappa.raw	<i>Fleiss' generalized kappa among multiple raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.</i>
------------------	--

Description

Fleiss' generalized kappa among multiple raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.

Usage

```
fleiss.kappa.raw(ratings, weights = "unweighted", categ.labels = NULL,
  conflev = 0.95, N = Inf)
```

Arguments

ratings	An $n \times r$ matrix / data frame of ratings where each column represents one rater and each row one subject.
---------	---

weights	is a mandatory parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix $q \times q$ where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter <code>categ.labels</code> . Otherwise, the program may not work.
categ.labels	An optional vector parameter containing the list of all possible ratings. It may be useful in case some of the possible ratings are not used by any rater, they will still be used when calculating agreement coefficients. The default value is NULL. In this case, only categories reported by the raters are used in the calculations.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A data list containing 3 objects: (1) a one-row data frame containing various statistics including the requested agreement coefficient, (2) the weight matrix used in the calculations if any, and (3) the categories used in the analysis. These could be categories reported by the raters, or those that were available to the raters whether they used them or not. The output data frame contains the following variables: "coeff.name" (coefficient name-here it will be "Fleiss' Kappa"), "pa" (the percent agreement), "pe" (the percent chance agreement), `coeff.val` (the agreement coefficient estimate-Fleiss' Kappa), "coeff.se" (the standard error), "conf.int" (Fleiss Kappa's confidence interval), "p.value" (Fleiss Kappa's p-value), "w.name" (the weights' identification).

References

Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. John Wiley & Sons.

gwet.ac1.dist	<i>Gwet's AC1/AC2 agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.</i>
---------------	---

Description

Gwet's AC1/AC2 agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.

Usage

```
gwet.ac1.dist(ratings, weights = "unweighted", categ = NULL,
             conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxq matrix / data frame containing the distribution of raters by subject and category. Each cell (i,k) contains the number of raters who classified subject i into category k .
weights	is an optional parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter <code>categ</code> . Otherwise, only the categories reported will be used.
categ	An optional parameter representing all categories available to raters during the experiment. This parameter may be useful if some categories were not used by any rater in spite of being available to the raters.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A vector containing the following information: `pa`(the percent agreement), `pe`(the percent chance agreement), `coeff`(Gwet's AC1 or AC2 depending on whether weights are used or not), `stderr`(the standard error of Gwet's coefficient), `conf.int`(the confidence interval of Gwet's coefficient), `p.value`(the p-value of Gwet's coefficient), `coeff.name` (AC1/AC2).

Source

Gwet, K. L. (2008). "Computing inter-rater reliability and its variance in the presence of high agreement," *British Journal of Mathematical and Statistical Psychology*, 61, 29-48.

gwet.ac1.raw	<i>Gwet's AC1/AC2 agreement coefficient among multiple raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.</i>
--------------	---

Description

Gwet's AC1/AC2 agreement coefficient among multiple raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.

Usage

```
gwet.ac1.raw(ratings, weights = "unweighted", categ.labels = NULL,
             conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxr matrix / data frame of ratings where each column represents one rater and each row one subject.
weights	is a mandatory parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter categ.labels. Otherwise, the program may not work.
categ.labels	An optional vector parameter containing the list of all possible ratings. It may be useful in case some of the possible ratings are not used by any rater, they will still be used when calculating agreement coefficients. The default value is NULL. In this case, only categories reported by the raters are used in the calculations.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A data list containing 3 objects: (1) a one-row data frame containing various statistics including the requested agreement coefficient, (2) the weight matrix used in the calculations if any, and (3) the categories used in the analysis. These could be categories reported by the raters, or those that were available to the raters whether they used them or not. The output data frame contains the following variables: "coeff.name" (coefficient name), "pa" (the percent agreement), "pe" (the percent chance agreement), "coeff.val" (the agreement coefficient estimate-AC1 or AC2), "coeff.se" (the standard error), "conf.int" (AC1/AC2 confidence interval), "p.value" (Gwet AC1/AC2 p-value), "w.name" (the weights' identification).

References

Gwet, K. L. (2008). "Computing inter-rater reliability and its variance in the presence of high agreement." *British Journal of Mathematical and Statistical Psychology*, 61, 29-48.

 gwet.ac1.table

Gwet's AC1/AC2 coefficient for 2 raters

Description

Gwet's AC1/AC2 coefficient for 2 raters

Usage

```
gwet.ac1.table(ratings, weights = identity.weights(1:ncol(ratings)),
  conflev = 0.95, N = Inf)
```

Arguments

ratings	A square table of ratings (assume no missing ratings).
weights	An optional matrix that contains the weights used in the weighted analysis. By default, this parameter contains the identity weight matrix, which leads to the unweighted analysis.
conflev	An optional parameter that specifies the confidence level used for constructing confidence intervals. By default the function assumes the standard value of 95%.
N	An optional parameter representing the finite population size if any. It is used to perform the finite population correction to the standard error. It's default value is infinity.

Value

A data frame containing the following 5 variables: coeff.name coeff.val coeff.se coeff.ci coeff.pval.

identity.weights	<i>Function for computing the Identity Weights</i>
------------------	--

Description

Function for computing the Identity Weights

Usage

```
identity.weights(categ)
```

Arguments

categ	A mandatory parameter representing the vector of all possible ratings.
-------	--

Value

A square matrix of identity weights to be used for calculating the unweighted coefficients.

kappa2.table	<i>Kappa coefficient for 2 raters</i>
--------------	---------------------------------------

Description

Kappa coefficient for 2 raters

Usage

```
kappa2.table(ratings, weights = identity.weights(1:ncol(ratings)),
             conflev = 0.95, N = Inf)
```

Arguments

ratings	A square table of ratings (assume no missing ratings).
weights	An optional matrix that contains the weights used in the weighted analysis.

Value

A data frame containing the following 5 variables: coeff.name coeff.val coeff.se coeff.ci coeff.pval.

krippen.alpha.dist	<i>Krippendorff's agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.</i>
--------------------	---

Description

Krippendorff's agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.

Usage

```
krippen.alpha.dist(ratings, weights = "unweighted", categ = NULL,
  conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxq matrix / data frame containing the distribution of raters by subject and category. Each cell (i,k) contains the number of raters who classified subject i into category k .
weights	is an optional parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter <code>categ</code> . Otherwise, only the categories reported will be used.
categ	An optional parameter representing all categories available to raters during the experiment. This parameter may be useful if some categories were not used by any rater in spite of being available to the raters.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A vector containing the following information: pa(the percent agreement),pe(the percent chance agreement),coeff(Krippendorff's alpha), stderr(the standard error of Krippendorff's coefficient),conf.int(the confidence interval of Krippendorff's alpha coefficient), p.value(the p-value of Krippendorff's alpha), coeff.name ("krippen alpha").

Source

Gwet, K. (2014). *Handbook of Inter-Rater Reliability: The Definitive Guide to Measuring the Extent of Agreement Among Multiple Raters*, 4th Edition. Advanced Analytics, LLC
 Krippendorff (1970). "Bivariate agreement coefficients for reliability of data," *Sociological Methodology*, 2, 139-150
 Krippendorff (1980). *Content analysis: An introduction to its methodology (2nd ed.)*, Newbury Park, CA: Sage.

krippen.alpha.raw	<i>Krippendorff's alpha coefficient for an arbitrary number of raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.</i>
-------------------	---

Description

Krippendorff's alpha coefficient for an arbitrary number of raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.

Usage

```
krippen.alpha.raw(ratings, weights = "unweighted", categ.labels = NULL,
  conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxr matrix / data frame of ratings where each column represents one rater and each row one subject.
weights	is a mandatory parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter categ.labels. Otherwise, the program may not work.
categ.labels	An optional vector parameter containing the list of all possible ratings. It may be useful in case some of the possible ratings are not used by any rater, they will still be used when calculating agreement coefficients. The default value is NULL. In this case, only categories reported by the raters are used in the calculations.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A data list containing 3 objects: (1) a one-row data frame containing various statistics including the requested agreement coefficient—in this case, Krippendorff's alpha, (2) the weight matrix used in the calculations if any, and (3) the vector of categories used in the analysis. These could be categories reported by the raters, or those that were available to the raters whether they used them

or not. The output data frame contains the following variables: "coeff.name" (coefficient name), "pa" (the percent agreement), "pe" (the percent chance agreement), coeff.val (Krippendorff's alpha estimate), "coeff.se" (standard error), conf.int" (Krippendorff alpha's confidence interval), "p.value" (Krippendorff alpha's p-value), "w.name" (the weights' identification).

References

Gwet, K. (2014). *Handbook of Inter-Rater Reliability: The Definitive Guide to Measuring the Extent of Agreement Among Multiple Raters*, 4th Edition. Advanced Analytics, LLC.

Krippendorff (1970). "Bivariate agreement coefficients for reliability of data." *Sociological Methodology*, 2, 139-150.

Krippendorff (1980). *Content analysis: An introduction to its methodology* (2nd ed.), New-bury Park, CA: Sage.

krippen2.table	<i>Krippendorff's Alpha coefficient for 2 raters</i>
----------------	--

Description

Krippendorff's Alpha coefficient for 2 raters

Usage

```
krippen2.table(ratings, weights = identity.weights(1:ncol(ratings)),
  conflev = 0.95, N = Inf)
```

Arguments

ratings	A square table of ratings (assume no missing ratings).
weights	An optional matrix that contains the weights used in the weighted analysis. By default, this parameter contains the identity weight matrix, which leads to the unweighted analysis.
conflev	An optional parameter that specifies the confidence level used for constructing confidence intervals. By default the function assumes the standard value of 95%.
N	An optional parameter representing the finite population size if any. It is used to perform the finite population correction to the standard error. It's default value is infinity.

Value

A data frame containing the following 5 variables: coeff.name coeff.val coeff.se coeff.ci coeff.pval.

 landis.koch

Dataset describing the Landis & Koch Benchmarking Scale

Description

This dataset contains information describing the Landis & Koch scale for benchmarking chance-corrected agreement coefficients such as Gwet AC1/AC2, Kappa and many others.

Usage

```
landis.koch
```

Format

Each row of this dataset describes an interval and the interpretation of the magnitude it represents.

lb.LK The interval lower bound

ub.LK The interval upper bound

interp.LK The interpretation

Source

Landis, J.R. & Koch G. (1977). The measurement of observer agreement for categorical data, *Biometrics*, 33, 159-174.

 landis.koch.bf

Computing Landis-Koch Benchmark Scale Membership Probabilities

Description

Computing Landis-Koch Benchmark Scale Membership Probabilities

Usage

```
landis.koch.bf(coeff, se, BenchDF = landis.koch)
```

Arguments

coeff	A mandatory parameter representing the estimated value of an agreement coefficient.
se	A mandatory parameter representing the agreement coefficient standard error.
BenchDF	An optional parameter that is a 3-column data frame containing the Landis & Koch's benchmark scale information. The 3 columns are the interval lower bound, upper bound, and their interpretation. The default value is a small file contained in the package and named <i>landis.koch.RData</i> , which describes the official Landis & Koch's scale intervals and their interpretation.

Value

A one-column matrix containing the membership probabilities (c.f. http://agreestat.com/research_papers/inter-rater%20reliability%20study%20design1.pdf)

linear.weights	<i>Function for computing the Linear Weights</i>
----------------	--

Description

Function for computing the Linear Weights

Usage

```
linear.weights(categ)
```

Arguments

categ A mandatory parameter representing the vector of all possible ratings.

Value

A square matrix of quadratic weights to be used for calculating the weighted coefficients.

ordinal.weights	<i>Function for computing the Ordinal Weights</i>
-----------------	---

Description

Function for computing the Ordinal Weights

Usage

```
ordinal.weights(categ)
```

Arguments

categ A mandatory parameter representing the vector of all possible ratings.

Value

A square matrix of quadratic weights to be used for calculating the weighted coefficients.

pa.coeff.dist	<i>Percent agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.</i>
---------------	--

Description

Percent agreement coefficient among multiple raters (2, 3, +) when the input dataset is the distribution of raters by subject and category.

Usage

```
pa.coeff.dist(ratings, weights = "unweighted", categ = NULL,
             conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxq matrix / data frame containing the distribution of raters by subject and category. Each cell (i,k) contains the number of raters who classified subject i into category k .
weights	is an optional parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter <code>categ</code> . Otherwise, only the categories reported will be used.
categ	An optional parameter representing all categories available to raters during the experiment. This parameter may be useful if some categories were not used by any rater in spite of being available to the raters.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A vector containing the following information: pa(the percent agreement),pe(the percent chance agreement),coeff(Brennan-Prediger coefficient), stderr(the standard error of Brennan-Prediger coefficient),conf.int(the p-value of Brennan-Prediger coefficient), p.value(the p-value of Brennan-Prediger coefficient),coeff.name ("Brennan-Prediger").

Source

Brennan, R.L., and Prediger, D. J. (1981). "Coefficient Kappa: some uses, misuses, and alternatives," *Educational and Psychological Measurement*, 41, 687-699.

pa.coeff.raw	<i>Percent agreement among multiple raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.</i>
--------------	--

Description

Percent agreement among multiple raters (2, 3, +) when the input data represent the raw ratings reported for each subject and each rater.

Usage

```
pa.coeff.raw(ratings, weights = "unweighted", categ.labels = NULL,
             conflev = 0.95, N = Inf)
```

Arguments

ratings	An nxr matrix / data frame of ratings where each column represents one rater and each row one subject.
weights	is a mandatory parameter that is either a string variable or a matrix. The string describes one of the predefined weights and must take one of the values ("quadratic", "ordinal", "linear", "radical", "ratio", "circular", "bipolar"). If this parameter is a matrix then it must be a square matrix qxq where q is the number of possible categories where a subject can be classified. If some of the q possible categories are not used, then it is strongly advised to specify the complete list of possible categories as a vector in parameter categ.labels. Otherwise, the program may not work.
categ.labels	An optional vector parameter containing the list of all possible ratings. It may be useful in case some of the possible ratings are not used by any rater, they will still be used when calculating agreement coefficients. The default value is NULL. In this case, only categories reported by the raters are used in the calculations.
conflev	An optional parameter representing the confidence level associated with the confidence interval. Its default value is 0.95.
N	An optional parameter representing the population size (if any). It may be used to perform the final population correction to the variance. Its default value is infinity.

Value

A data list containing 3 objects: (1) a one-row data frame containing the estimates, (2) the weight matrix used in the calculations, and (3) the categories used in the analysis. The data frame of estimates contains the following variables "coeff.name" (coefficient name), "pa" (the percent agreement), "pe" (percent chance-agreement-always equals 0), "coeff.val" (agreement coefficient = pa), "coeff.se" (the percent agreement standard error), "conf.int" (the percent agreement confidence interval), "p.value" (the percent agreement p-value), "w.name" (the weights' identification).

pa2.table *Percent Agreement coefficient for 2 raters*

Description

Percent Agreement coefficient for 2 raters

Usage

```
pa2.table(ratings, weights = identity.weights(1:ncol(ratings)),
          conflev = 0.95, N = Inf)
```

Arguments

ratings	A square table of ratings (assume no missing ratings).
weights	An optional matrix that contains the weights used in the weighted analysis. By default, this parameter contains the identity weight matrix, which leads to the unweighted analysis.
conflev	An optional parameter that specifies the confidence level used for constructing confidence intervals. By default the function assumes the standard value of 95%.
N	An optional parameter representing the finite population size if any. It is used to perform the finite population correction to the standard error. It's default value is infinity.

Value

A data frame containing the following 5 variables: coeff.name coeff.val coeff.se coeff.ci coeff.pval.

quadratic.weights *Function for computing the Quadratic Weights*

Description

Function for computing the Quadratic Weights

Usage

```
quadratic.weights(categ)
```

Arguments

categ	A mandatory parameter representing the vector of all possible ratings.
-------	--

Value

A square matrix of quadratic weights to be used for calculating the weighted coefficients.

radical.weights	<i>Function for computing the Radical Weights</i>
-----------------	---

Description

Function for computing the Radical Weights

Usage

```
radical.weights(categ)
```

Arguments

categ A mandatory parameter representing the vector of all possible ratings.

Value

A square matrix of quadratic weights to be used for calculating the weighted coefficients.

ratio.weights	<i>Function for computing the Ratio Weights</i>
---------------	---

Description

Function for computing the Ratio Weights

Usage

```
ratio.weights(categ)
```

Arguments

categ A mandatory parameter representing the vector of all possible ratings.

Value

A square matrix of quadratic weights to be used for calculating the weighted coefficients.

scott2.table	<i>Scott's coefficient for 2 raters</i>
--------------	---

Description

Scott's coefficient for 2 raters

Usage

```
scott2.table(ratings, weights = identity.weights(1:ncol(ratings)),
  conflev = 0.95, N = Inf)
```

Arguments

ratings	A square table of ratings (assume no missing ratings).
weights	An optional matrix that contains the weights used in the weighted analysis. By default, this parameter contains the identity weight matrix, which leads to the unweighted analysis.
conflev	An optional parameter that specifies the confidence level used for constructing confidence intervals. By default the function assumes the standard value of 95%.
N	An optional parameter representing the finite population size if any. It is used to perform the finite population correction to the standard error. It's default value is infinity.

Value

A data frame containing the following 5 variables: coeff.name coeff.val coeff.se coeff.ci coeff.pval.

trim	<i>An r function for trimming leading and trailing blanks</i>
------	---

Description

An r function for trimming leading and trailing blanks

Usage

```
trim(x)
```

Arguments

x	is a string variable.
---	-----------------------

Value

A string variable where leading and trailing blanks are trimmed.

Index

*Topic **datasets**

- altman, [2](#)
 - cac.ben.gerry, [6](#)
 - cac.dist.g1g2, [7](#)
 - cac.dist4cat, [8](#)
 - cac.raw.g1g2, [8](#)
 - cac.raw.gender, [9](#)
 - cac.raw4raters, [9](#)
 - cac.raw5obser, [10](#)
 - cont3x3abstractors, [12](#)
 - cont4x4diagnosis, [13](#)
 - distrib.6raters, [13](#)
 - fleiss, [14](#)
 - landis.koch, [24](#)
- altman, [2](#)
- altman.bf, [3](#)
- bipolar.weights, [3](#)
- bp.coeff.dist, [4](#)
- bp.coeff.raw, [5](#)
- bp2.table, [6](#)
- cac.ben.gerry, [6](#)
- cac.dist.g1g2, [7](#)
- cac.dist4cat, [8](#)
- cac.raw.g1g2, [8](#)
- cac.raw.gender, [9](#)
- cac.raw4raters, [9](#)
- cac.raw5obser, [10](#)
- circular.weights, [11](#)
- conger.kappa.raw, [11](#)
- cont3x3abstractors, [12](#)
- cont4x4diagnosis, [13](#)
- distrib.6raters, [13](#)
- fleiss, [14](#)
- fleiss.bf, [15](#)
- fleiss.kappa.dist, [15](#)
- fleiss.kappa.raw, [16](#)
- gwet.ac1.dist, [17](#)
- gwet.ac1.raw, [18](#)
- gwet.ac1.table, [19](#)
- identity.weights, [20](#)
- kappa2.table, [20](#)
- krippen.alpha.dist, [21](#)
- krippen.alpha.raw, [22](#)
- krippen2.table, [23](#)
- landis.koch, [24](#)
- landis.koch.bf, [24](#)
- linear.weights, [25](#)
- ordinal.weights, [25](#)
- pa.coeff.dist, [26](#)
- pa.coeff.raw, [27](#)
- pa2.table, [28](#)
- quadratic.weights, [28](#)
- radical.weights, [29](#)
- ratio.weights, [29](#)
- scott2.table, [30](#)
- trim, [30](#)