Package 'dunn.test'

April 12, 2024

Version 1.3.6

Date 2024-4-12

Title Dunn's Test of Multiple Comparisons Using Rank Sums

Author Alexis Dinno

Maintainer Alexis Dinno <alexis.dinno@pdx.edu>

Description Computes Dunn's test (1964) for stochastic dominance and reports the results among multiple pairwise comparisons after a Kruskal-Wallis test for 0th-order stochastic dominance among k groups (Kruskal and Wallis, 1952). 'dunn.test' makes k(k-1)/2 multiple pairwise comparisons based on Dunn's z-test-statistic approximations to the actual rank statistics. The null hypothesis for each pairwise comparison is that the probability of observing a randomly selected value from the first group that is larger than a randomly selected value from the second group equals one half; this null hypothesis corresponds to that of the Wilcoxon-Mann-Whitney rank-sum test. Like the rank-sum test, if the data can be assumed to be continuous, and the distributions are assumed identical except for a difference in location, Dunn's test may be understood as a test for median difference and for mean difference. 'dunn.test' accounts for tied ranks.

License GPL-2

LazyData true

Encoding UTF-8

NeedsCompilation no

Repository CRAN

Date/Publication 2024-04-12 18:30:02 UTC

R topics documented:

dunn.test *Dunn's Test*

Description

Performs Dunn's test of multiple comparisons using rank sums

Usage

```
dunn.test  (x, g=NA, method=p.adjustment.methods, kw=TRUE, label=TRUE,
      wrap=FALSE, table=TRUE, list=FALSE, rmc=FALSE, alpha=0.05, altp=FALSE)

p.adjustment.methods
# c("none", "bonferroni", "sidak", "holm", "hs", "hochberg", "bh", "by")
```

Arguments

x	a numeric vector, or a list of numeric vectors. Missing values are ignored. If the former, then groups *must* be specified using g.
g	a factor variable, numeric vector, or character vector indicating group. Missing values are ignored.
method	adjusts the *p*-value for multiple comparisons using the Bonferroni, Šidák, Holm, Holm-Šidák, Hochberg, Benjamini-Hochberg, or Benjamini-Yekutieli adjustment (see Details). The default is no adjustment for multiple comparisons.
kw	if 'TRUE' then the results of the Kruskal-Wallis test are reported.
label	if 'TRUE' then the factor labels are used in the output table.
wrap	does not break up tables to maintain nicely formatted output. If 'FALSE' then output of large tables is broken up across multiple pages.
table	outputs results of Dunn's test in a table format, as qualified by the label and wrap options.
list	outputs results of Dunn's test in a list format.
rmc	if 'TRUE' then the reported test statistics and table are based on row minus column, rather than the default column minus row (i.e. the signs of the test statistic are flipped).
alpha	the nominal level of significance used in the step-up/step-down multiple comparisons procedures (Holm, Holm-Šidák, Hochberg, Benjamini-Hochberg, and Benjamini-Yekutieli).
altp	if 'TRUE' then express *p*-values in alternative format. The default is to express *p*-value = P($Z \geq \|z\|$), and reject Ho if $p \leq \alpha/2$. When the altp option is used, *p*-values are instead expressed as *p*-value = P($\|Z\| \geq \|z\|$), and reject Ho if $p \leq \alpha$. **These two expressions give identical test results.** Use of altp is therefore merely a semantic choice.

Details

dunn.test computes Dunn's test (1964) for stochastic dominance and reports the results among multiple pairwise comparisons after a Kruskal-Wallis test for stochastic dominance among k groups (Kruskal and Wallis, 1952). The interpretation of stochastic dominance requires an assumption that the CDF of one group does not cross the CDF of the other. dunn.test makes $m = k(k\text{-}1)/2$ multiple pairwise comparisons based on Dunn's z-test-statistic approximations to the actual rank statistics. The null hypothesis for each pairwise comparison is that the probability of observing a randomly selected value from the first group that is larger than a randomly selected value from the second group equals one half; this null hypothesis corresponds to that of the Wilcoxon-Mann-Whitney rank-sum test. Like the rank-sum test, if the data can be assumed to be continuous, and the distributions are assumed identical except for a difference in location, Dunn's test may be understood as a test for median difference. dunn.test accounts for tied ranks.

dunn.test outputs both z-test-statistics for each pairwise comparison and the p-value = P($Z \geq |z|$) for each. Reject Ho based on $p \leq \alpha/2$ (and in combination with p-value ordering for stepwise method options). If you prefer to work with p-values expressed as p-value = P($|Z| \geq |z|$) use the altp=TRUE option, and reject Ho based on $p \leq \alpha$ (and in combination with p-value ordering for stepwise method options). These are exactly equivalent rejection decisions.

Several options are available to adjust p-values for multiple comparisons, including methods to control the family-wise error rate (FWER) and methods to control the false discovery rate (FDR):

'none' no adjustment is made. Those comparisons rejected without adjustment at the α level (two-sided test) are starred in the output table, and starred in the list when using the list=TRUE option.

'bonferroni' the FWER is controlled using Dunn's (1961) Bonferroni adjustment, and adjusted p-values = max(1, pm). Those comparisons rejected with the Bonferroni adjustment at the α level (two-sided test) are starred in the output table, and starred in the list when using the list=TRUE option.

'sidak' the FWER is controlled using Šidák's (1967) adjustment, and adjusted p-values = max(1, 1 - (1 - p)^m). Those comparisons rejected with the Šidák adjustment at the α level (two-sided test) are starred in the output table, and starred in the list when using the list=TRUE option.

'holm' the FWER controlled using Holm's (1979) progressive step-up procedure to relax control on subsequent tests. p values are ordered from smallest to largest, and adjusted p-values = max[1, $p(m+1\text{-}i)$], where i indexes the ordering. All tests after and including the first test to not be rejected are also not rejected.

'hs' the FWER is controlled using the Holm-Šidák adjustment (Holm, 1979): another progressive step-up procedure but assuming dependence between tests. p values are ordered from smallest to largest, and adjusted p-values = max[1, 1 - (1 - p)^(m+1-i)], where i indexes the ordering. All tests after and including the first test to not be rejected are also not rejected.

'hochberg' the FWER is controlled using Hochberg's (1988) progressive step-down procedure to increase control on successive tests. p values are ordered from largest–

smallest, and adjusted p-values = max[1, p*i], where i indexes the ordering. All tests after and including the first to be rejected are also rejected.

'bh' the FDR is controlled using the Benjamini-Hochberg adjustment (1995), a step-down procedure appropriate to independent tests or tests that are positively dependent. p-values are ordered from largest to smallest, and adjusted p-values = max[1, $pm/(m+1-i)$], where i indexes the ordering. All tests after and including the first to be rejected are also rejected.

'by' the FDR is controlled using the Benjamini-Yekutieli adjustment (2011), a step-down procedure appropriate to dependent tests. p-values are ordered from largest to smallest, and adjusted p-values = max[1, $pmC/(m+1-i)$], where i indexes the ordering, and the constant $C = 1 + 1/2 + \ldots + 1/m$. All tests after and including the first to be rejected are also rejected.

Because the sequential step-up/step-down tests rejection decisions depend on both the p-values and their ordering, those tests rejected using `"holm"`, `"hs"`, `"hochberg"`, `"bh"`, or `"by"` at the indicated α level are starred in the output table, and starred in the list when using the `list=TRUE` option.

Value

`dunn.test` returns:

`chi2`	a scalar of the Kruskal-Wallis test statistic adjusted for ties.
`Z`	a vector of all m of Dunn z test statistics.
`P`	a vector of p-values corresponding to `Z`. –OR–
`altP`	a vector of p-values corresponding to `Z` when using the `altp=TRUE` option.
`P.adjust`	a vector of p-values corresponding to `Z`, but adjusted for multiple comparisons as per `method` (`P = P.adjust` if `method="none"`). –OR–
`altP.adjust`	a vector of p-values corresponding to `Z`, but adjusted for multiple comparisons as per `method` (`P = P.adjust` if `method="none"`) when using the `altp=TRUE` option.
`comparisons`	a vector of strings labeling each pairwise comparison, as qualified by the `rmc` option, using either the variable values, or the factor labels or (or factor values if unlabeled). These labels match the corresponding position in the `Z`, `P`, and `P.adjust` vectors.

Author(s)

Alexis Dinno (<alexis.dinno@pdx.edu>)

Please contact me with any questions, bug reports or suggestions for improvement. Fixing bugs will be facilitated by sending along:

[1] a copy of the data (de-labeled or anonymized is fine),
[2] a copy of the command syntax used, and
[3] a copy of the exact output of the command.

References

Benjamini, Y. and Hochberg, Y. (1995) Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*. **57**, 289–300. <doi:10.1111/j.25176161.1995.tb02031.x>.

Benjamini, Y. and Yekutieli, D. (2001) The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*. **29**, 1165–1188. <doi:10.1214/aos/1013699998>.

Dunn, O. J. (1961) Multiple comparisons among means. *Journal of the American Statistical Association*. **56**, 52–64. <doi:10.1080/01621459.1961.10482090>.

Dunn, O. J. (1964) Multiple comparisons using rank sums. *Technometrics*. **6**, 241–252. <doi:10.1080/00401706.1964.10490181>.

Hochberg, Y. (1988) A sharper Bonferroni procedure for multiple tests of significance. *Biometrika*. **75**, 800–802. <doi:10.1093/biomet/75.4.800>.

Holm, S. (1979) A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*. **6**, 65–70.

Kruskal, W. H. and Wallis, A. (1952) Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association*. **47**, 583–621. <doi:10.1080/01621459.1952.10483441>.

Šidák, Z. (1967) Rectangular confidence regions for the means of multivariate normal distributions. *Journal of the American Statistical Association*. **62**, 626–633. <doi:10.1080/01621459.1967.10482935>.

Examples

```
## Example cribbed and modified from the kruskal.test documentation
## Hollander & Wolfe (1973), 116.
## Mucociliary efficiency from the rate of removal of dust in normal
##  subjects, subjects with obstructive airway disease, and subjects
##  with asbestosis.
x <- c(2.9, 3.0, 2.5, 2.6, 3.2) # normal subjects
y <- c(3.8, 2.7, 4.0, 2.4)      # with obstructive airway disease
z <- c(2.8, 3.4, 3.7, 2.2, 2.0) # with asbestosis
dunn.test(x=list(x,y,z))

x <- c(x, y, z)
g <- factor(rep(1:3, c(5, 4, 5)),
            labels = c("Normal",
                       "COPD",
                       "Asbestosis"))
dunn.test(x, g)

## Example based on home care data from Dunn (1964)
data(homecare)
attach(homecare)
dunn.test(occupation, eligibility, method="hs", list=TRUE)

## Air quality data set illustrates differences in different
## multiple comparisons adjustments
attach(airquality)
dunn.test(Ozone, Month, kw=FALSE, method="bonferroni")
dunn.test(Ozone, Month, kw=FALSE, method="hs")
```

```
dunn.test(Ozone, Month, kw=FALSE, method="bh")
detach(airquality)
```

`homecare` *Occupation and Home Care Eligibility*

Description

Occupation and home care eligibility for 383 patients medically elgibile for home care.

Usage

```
homecare
```

Format

A data frame containing two variables, `occupation` and `eligibility`, over 383 observations.

Source

Dunn, O. J. (1964) Multiple comparisons using rank sums. *Technometrics*, **6(3)**, 241–252s. Table 1. <doi:10.1080/00401706.1964.10490181>.

References

The data have been adapted from a study of a group of patients entering the Los Angeles County General Hospital during the years 1959-61.

Index