Package: fcrci (via r-universe)

November 13, 2026

Type Package

Title FDR-Consistent Confidence Intervals for Multiple Comparisons

Version 0.1.0

Description Computes confidence intervals that are consistent with false discovery rate (FDR) adjusted p-values, using the false coverage-statement rate (FCR) adjustment of Benjamini and Yekutieli (2005) <doi:10.1198/016214504000001907>. When p-values are adjusted with the Benjamini-Hochberg or Benjamini-Yekutieli procedures, the intervals exclude the null value exactly for the hypotheses that are rejected, so intervals and adjusted p-values agree by construction. The first release supports contrasts produced by the 'emmeans' package.

License GPL (>= 3)

URL https://github.com/jiarmendano/fcrci

BugReports https://github.com/jiarmendano/fcrci/issues

Encoding UTF-8

Language en-US

Depends R (>= 4.1.0)

Imports stats

Suggests emmeans (>= 1.10.0), knitr, rmarkdown, testthat (>= 3.0.0)

VignetteBuilder knitr

Config/testthat/edition 3

Roxygen list(markdown = TRUE)

Config/roxygen2/version 8.1.0

Repository https://jiarmendano.r-universe.dev

Date/Publication 2026-10-27 18:09:31 UTC

RemoteUrl https://github.com/jiarmendano/fcrci

RemoteRef HEAD

RemoteSha 9af1dcf53302fcd49e3901e0f815b615f5953c56

Contents

confint_fcr 2
confint_fcr.emmGrid 3
fcr_info 5
fcr_level 6

Index **8**

`confint_fcr` *FDR-consistent (FCR-adjusted) confidence intervals*

Description

Computes confidence intervals that agree with false discovery rate (FDR) adjusted p-values, using the false coverage-statement rate (FCR) adjustment of Benjamini and Yekutieli (2005). Within each family of tests, an interval excludes the null value if and only if the corresponding FDR-adjusted p-value is at most `1 - level`.

Usage

```
confint_fcr(object, ...)

## Default S3 method:
confint_fcr(object, ...)
```

Arguments

`object`	An object for which FCR-adjusted intervals are wanted.
`...`	Further arguments passed to methods.

Details

`confint_fcr()` is a generic function. The current release provides methods for `emmGrid` and `emm_list` objects created by the 'emmeans' package, typically the result of `emmeans::contrast()` or `pairs()`.

Value

See the documentation of the specific method, for example `confint_fcr.emmGrid()`.

References

Benjamini, Y. and Yekutieli, D. (2005). False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469), 71-81. doi:10.1198/016214504000001907

See Also

`confint_fcr.emmGrid()`, `fcr_level()`, `fcr_info()`.

`confint_fcr.emmGrid` *FCR-adjusted confidence intervals for 'emmeans' contrasts*

Description

Returns the usual 'emmeans' summary table (estimates, confidence intervals, test statistics and p-values) in which the p-values are FDR-adjusted and the confidence intervals are FCR-adjusted, so that both lead to the same conclusions.

Usage

```
## S3 method for class 'emmGrid'
confint_fcr(object, level = NULL, method = NULL, ...)

## S3 method for class 'emm_list'
confint_fcr(object, level = NULL, method = NULL, ..., which = NULL)
```

Arguments

`object`	An `emmGrid` object, usually the result of `emmeans::contrast()` or `pairs()`. For `emm_list` objects (for example from `emmeans(fit, pairwise ~ x)`), the `contrasts` component is used by default.
`level`	Confidence level. The FDR and FCR level is `q = 1 - level`. If `NULL` (the default), the level stored in `object` is used, which is 0.95 unless it was changed with `update()`.
`method`	`"BH"` (Benjamini-Hochberg, identical to `adjust = "fdr"` in 'emmeans') or `"BY"` (Benjamini-Yekutieli). If `NULL` (the default), the method is taken from the `adjust` setting stored in `object`, which must then be `"fdr"`, `"BH"` or `"BY"`.
`...`	Further arguments passed to `emmeans::summary.emmGrid()`, such as `type = "response"`, `by`, `side`, `null` or `df`. The arguments `adjust`, `cross.adjust`, `infer` and `delta` (or abbreviations of them) are not allowed because they would change the family or break the agreement between intervals and p-values.
`which`	For `emm_list` objects, the name or index of the component to use. Defaults to `"contrasts"` when that component exists.

Value

A `summary_emm` data frame, printed in the usual 'emmeans' style, with FCR-adjusted confidence limits and FDR-adjusted p-values. The number of tests, the number of rejections and the confidence level used in each family are stored as an attribute and can be retrieved with `fcr_info()`.

Why this is needed

In 'emmeans', the `adjust` argument accepts FDR methods such as `"fdr"` (`"BH"`) or `"BY"` for p-values, but confidence intervals cannot be FDR-adjusted, so `confint()` switches to a Bonferroni

adjustment (noted only in the footer of the output). As a result, a contrast can be significant according to its FDR-adjusted p-value while its (Bonferroni-adjusted) confidence interval includes the null value. `confint_fcr()` replaces those intervals with the FCR-adjusted intervals of Benjamini and Yekutieli (2005), which are the interval counterpart of the Benjamini-Hochberg procedure.

Method

For each family of m tests, the FDR procedure rejects R hypotheses at level $q = 1 - \text{level}$. Marginal intervals are then computed at confidence level $1 - Rq/m$ (or $1 - Rq/(m\sum_{i=1}^{m} 1/i)$ for `method = "BY"`); see `fcr_level()`. When $R = 0$, the level is computed with $R = 1$. Benjamini and Yekutieli (2005) construct no intervals when nothing is selected; using $R = 1$ is a convention of this package that gives Bonferroni intervals (at level q for `"BH"`, at $q/\sum 1/i$ for `"BY"`), none of which excludes the null value.

The intervals themselves are computed by 'emmeans' (`summary(..., adjust = "none", level = ...)`), so degrees of freedom, back-transformation (for example ratios or odds ratios with `type = "response"`), non-zero `null` values and one-sided intervals (`side`) are handled exactly as in 'emmeans'.

Families are defined as in 'emmeans': each group of the `by` variables is a separate family, both for the p-value adjustment and for the interval adjustment. Use `by = NULL` to treat all contrasts as a single family.

What is and is not guaranteed

The agreement between intervals and adjusted p-values holds by construction (up to floating-point ties at the threshold). Benjamini and Yekutieli (2005) prove that the `"BH"` version controls the FCR, the expected proportion of non-covering intervals among the *selected* (significant) parameters, at level q when the test statistics are independent, and that the `"BY"` version controls it under any dependence structure. Contrasts estimated from the same model are usually correlated, so `method = "BY"` is the conservative choice when a formal guarantee is required. Intervals are also reported for the non-selected contrasts so that the table is complete, but the FCR criterion makes no coverage statement about them.

Equivalence tests (`delta`) and cross-family adjustments (`cross.adjust`) are not supported, whether they are passed as arguments or stored in `object` with `update()`.

References

Benjamini, Y. and Yekutieli, D. (2005). False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469), 71-81. doi:10.1198/016214504000001907

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*, 57(1), 289-300. doi:10.1111/j.25176161.1995.tb02031.x

See Also

`fcr_level()`, `fcr_info()`, `emmeans::summary.emmGrid()`.

Examples

```
library(emmeans)

fit <- lm(weight ~ feed, data = chickwts)
prs <- pairs(emmeans(fit, ~ feed), adjust = "fdr")

# FDR-adjusted p-values, but Bonferroni-adjusted intervals
confint(prs)

# FDR-adjusted p-values and FCR-adjusted intervals that agree
confint_fcr(prs)

# Ratios on the response scale, one family per level of 'percent'
fit2 <- lm(log(conc) ~ source * factor(percent), data = pigs)
prs2 <- pairs(emmeans(fit2, ~ source | percent), adjust = "fdr")
res <- confint_fcr(prs2, type = "response")
res
fcr_info(res)
```

---

fcr_info *Details of the FCR adjustment*

Description

Extracts, for each family of tests, the number of tests (`m`), the number of rejected hypotheses (`R`) and the confidence level used for the FCR-adjusted intervals.

Usage

```
fcr_info(x)
```

Arguments

`x`	A result of `confint_fcr()`.

Value

A data frame with columns `family`, `m`, `R` and `level`.

Examples

```
library(emmeans)
fit <- lm(weight ~ feed, data = chickwts)
res <- confint_fcr(pairs(emmeans(fit, ~ feed), adjust = "fdr"))
fcr_info(res)
```

`fcr_level` *Confidence level for FCR-adjusted confidence intervals*

Description

Computes the confidence level that makes marginal confidence intervals consistent with Benjamini-Hochberg (BH) or Benjamini-Yekutieli (BY) adjusted p-values, following the false coverage-statement rate (FCR) adjustment of Benjamini and Yekutieli (2005).

Usage

```
fcr_level(p, q = 0.05, method = c("BH", "BY"))
```

Arguments

`p`	Numeric vector of unadjusted p-values forming one family of tests. Missing values are ignored.
`q`	FCR (and FDR) level, a single number strictly between 0 and 1. Usually `1 - level`, for example 0.05 for 95% intervals.
`method`	Either `"BH"` (Benjamini-Hochberg selection; `"fdr"`, as in `stats::p.adjust()` and 'emmeans', is accepted as a synonym) or `"BY"` (Benjamini-Yekutieli selection, valid under arbitrary dependence).

Details

Let m be the number of non-missing p-values in the family and R the number of hypotheses rejected at FDR level q. The FCR-adjusted confidence level is

$$1 - \frac{R\,q}{m\,c_m},$$

where $c_m = 1$ for `method = "BH"` and $c_m = \sum_{i=1}^{m} 1/i$ for `method = "BY"`.

Because the step-up procedure rejects exactly the hypotheses with $p_i \le Rq/(mc_m)$, a marginal interval computed at this level excludes the null value if and only if the corresponding adjusted p-value is at most q, provided that the interval and the unadjusted p-value come from the same test statistic and degrees of freedom.

When no hypothesis is rejected ($R = 0$), the level is computed with $R = 1$, that is $1 - q/(mc_m)$: the Bonferroni level for `"BH"`, and the Bonferroni level at q/c_m for `"BY"`. Benjamini and Yekutieli (2005) construct no intervals when nothing is selected; this convention of the package keeps intervals and adjusted p-values consistent, because no interval then excludes the null value.

Value

A single number: the confidence level to use for the marginal intervals of the family.

References

Benjamini, Y. and Yekutieli, D. (2005). False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469), 71-81. doi:10.1198/016214504000001907

Jung, K., Friede, T. and Beissbarth, T. (2011). Reporting FDR analogous confidence intervals for the log fold change of differentially expressed genes. *BMC Bioinformatics*, 12, 288. doi:10.1186/1471210512288

See Also

`confint_fcr()` to obtain the intervals for 'emmeans' contrasts.

Examples

```
p <- c(0.0004, 0.003, 0.012, 0.04, 0.21, 0.67)
lev <- fcr_level(p, q = 0.05)
lev

# The level is consistent with the BH-adjusted p-values:
# an interval at level `lev` excludes the null exactly when p <= 1 - lev.
p.adjust(p, method = "BH") <= 0.05
p <= 1 - lev

# Valid under arbitrary dependence, at the cost of wider intervals
fcr_level(p, q = 0.05, method = "BY")
```

Index

confint_fcr, 2
confint_fcr(), 5, 7
confint_fcr.emm_list
        (confint_fcr.emmGrid), 3
confint_fcr.emmGrid, 3
confint_fcr.emmGrid(), 2

emmeans::contrast(), 2, 3
emmeans::summary.emmGrid(), 3, 4

fcr_info, 5
fcr_info(), 2–4
fcr_level, 6
fcr_level(), 2, 4

stats::p.adjust(), 6