Package 'eClosure'

October 22, 2026

Type Package

Title Methods Based on the e-Closure Principle

Version 0.9.7

Maintainer Jelle Goeman <j.j.goeman@lumc.nl>

Description
Implements several methods for False Discovery Rate control based on the e-Closure Principle, in particular the Closed Benjamini-Hochberg, Closed e-Benjamini-Hochberg, Closed Su and Closed Benjamini-Yekutieli procedures.

License GPL-3

Imports Rcpp

LinkingTo Rcpp

NeedsCompilation yes

Encoding UTF-8

Suggests testthat (>= 3.0.0)

Config/testthat/edition 3

Config/roxygen2/version 8.1.1

Author Jelle Goeman [aut, cre]

Repository CRAN

Date/Publication 2026-10-21 23:20:02 UTC

Contents

closedBH 2
closedBY 3
closedeBH 5
closedSu 7
eClosure 9

Index **10**

`closedBH` *Closed Benjamini-Hochberg procedure for FDR control*

Description

Applies the e-Closure version of the Benjamini-Hochberg (BH) procedure. `closedBH()` returns the number of rejections at a given level, while `closedBH.adjust()` returns adjusted p-values, one per hypothesis.

Usage

```
closedBH(p, alpha = 0.05)

closedBH.adjust(p, cap = TRUE)
```

Arguments

`p`	Numeric vector of p-values, one per hypothesis. Values must lie in $[0, 1]$.
`alpha`	Numeric scalar in $[0, 1]$. The target FDR level. Defaults to `0.05`.
`cap`	Logical. If `TRUE` (the default), each adjusted p-value is raised to be at least its own raw p-value, so that adjustment never makes a p-value smaller. If `FALSE`, the raw adjusted values are returned.

Details

`closedBH()` returns the size r of the closedBH-significant set. The r hypotheses with the smallest p-values form one such set. This gives the maximum number of hypotheses that can be reported while maintaining FDR control.

`closedBH.adjust()` returns the adjusted p-value of each hypothesis: the smallest level α at which that hypothesis is among the rejections of `closedBH()` (capped at the p-value itself by default (`cap=TRUE`). Because the rejection set is always the hypotheses with the smallest p-values, a hypothesis of sorted rank k is rejected at level α exactly when `closedBH()` returns at least k; the adjusted p-value is the smallest such α. Consequently `sum(closedBH.adjust(p, cap = FALSE) <= alpha)` equals `closedBH(p, alpha = alpha)`.

All of the heavy lifting (the breadth-first traversal over set sizes and the built-in `next.r` recursion) is performed in C++; these wrappers only sort out the user-facing arguments and forward them. `closedBH()` runs in $O(m + r^2)$ time after sorting; `closedBH.adjust()` runs in $O(m^2)$ time and $O(m)$ memory.

Value

`closedBH()` returns a single non-negative integer r: the r hypotheses with the smallest p-values form a valid rejection set, and a value of `0` means no non-empty set can be rejected.

`closedBH.adjust()` returns a numeric vector of adjusted p-values, in the same order as the input `p`, each lying in $[0, 1]$.

References

Goeman, J. (2026). A Uniform Improvement of the Benjamini-Hochberg Procedure via e-Closure. arXiv preprint arXiv:2606.01854.

Benjamini, Y., & 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.

Xu, Z., Solari, A., Fischer, L., de Heide, R., Ramdas, A., & Goeman, J. (2025). Bringing closure to false discovery rate control: A general principle for multiple testing. arXiv preprint arXiv:2509.02517.

See Also

`closedBY()` for the analogous Benjamini-Yekutieli procedure. `closedSu()` for the analogous Su procedure. `closedeBH()` for the analogous procedure based on e-values. `p.adjust()` for standard non-simultaneous multiple testing corrections.

Examples

```
p <- c(
  0.0001, 0.0004, 0.0019, 0.0095, 0.0201,
  0.0278, 0.0298, 0.0344, 0.0459, 0.3240,
  0.4262, 0.5719, 0.6528, 0.7590, 1.0000
)
closedBH(p, alpha = 0.05)
closedBH(p, alpha = 0.10)

padj <- closedBH.adjust(p)
sum(padj <= 0.05)            # matches closedBH(p, alpha = 0.05)
```

`closedBY` *Closed Benjamini-Yekutieli procedure for simultaneous FDR control*

Description

Applies the closed testing version of the Benjamini-Yekutieli (BY) procedure. The standard BY procedure controls the false discovery rate (FDR) at level α under arbitrary dependence but only provides a single set of rejections. The closed BY procedure provides **simultaneous FDR control**: for every set of hypotheses, it determines whether that set can be reported as discoveries while maintaining FDR control at level α, regardless of which other sets are inspected.

Usage

```
closedBY(p, set = NULL, alpha = 0.05, approximate = FALSE)
```

Arguments

`p`	Numeric vector of p-values, one per hypothesis. Values must lie in $[0, 1]$; exact zeros are replaced internally by a small positive constant to avoid numerical issues.
`set`	Optional subsetting vector for `p` (logical, index, or negative index), indicating which hypotheses belong to the set to be checked for closedBY significance. If `NULL` (the default), the function instead returns the size of the largest closedBY-significant set.
`alpha`	Numeric scalar in $[0, 1]$. The target FDR level. Defaults to `0.05`.
`approximate`	Logical. If `FALSE` (the default), uses an exact algorithm that is guaranteed to find the largest closedBY-significant set. If `TRUE`, uses a faster approximate algorithm based on bisection that may occasionally return a smaller set. The approximate method is recommended for exploratory analyses or large inputs where computation time is a concern.

Details

The closed BY procedure is based on a local e-value for every intersection hypothesis. A set R of hypotheses is closedBY-significant — and therefore a valid simultaneous rejection — if and only if, for every subset $S \subseteq [m]$, the local e-value exceeds $|S \cap R|/(|R|\alpha)$. This guarantees post-hoc FDR control: you may report any closedBY-significant set as your discovery set without inflating the FDR above α, even if the choice of set was data-driven.

The function has two modes:

- **Set-checking mode** (when `set` is supplied): Returns `TRUE` if the specified set is closedBY-significant (i.e., can be reported as a valid simultaneous rejection at level α), and `FALSE` otherwise.
- **Discovery mode** (when `set = NULL`): Returns the size r of the largest closedBY-significant set. The r hypotheses with the smallest p-values always form one such set. This gives the maximum number of hypotheses that can be reported while maintaining simultaneous FDR control. In particular, the set R consisting of the r smallest p-values is closedBY-significant.

Note that closedBY significance is not a monotone property: a set of size r being closedBY-significant does not imply that all smaller sets are as well. The exact algorithm therefore checks all set sizes, while the approximate algorithm (`approximate = TRUE`) uses a faster bisection strategy that may occasionally underestimate the largest significant set.

Value

- If `set` is supplied: a single logical value. `TRUE` indicates that the specified set is closedBY-significant and can be reported as a simultaneous rejection at FDR level α. `FALSE` indicates it cannot.
- If `set = NULL`: a single non-negative integer r. The r hypotheses with the smallest p-values form a valid simultaneous rejection set. A return value of `0` means no non-empty set can be rejected.

References

Benjamini, Y., & Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4), 1165–1188.

Goeman, J. J., & Solari, A. (2011). Multiple testing for exploratory research. *Statistical Science*, 26(4), 584–597.

Xu, Z., Solari, A., Fischer, L., de Heide, R., Ramdas, A., & Goeman, J. (2025). Bringing closure to false discovery rate control: A general principle for multiple testing. arXiv preprint arXiv:2509.02517.

See Also

`closedBH()`, `closedSu()`, `closedeBH()`, eClosure

Examples

```
set.seed(42)
# 20 null hypotheses (p ~ Uniform(0,1)) and 10 non-nulls (p ~ Beta(0.1, 1), smaller on average)
p <- c(runif(20), rbeta(10, 0.1, 1))

# --- Discovery mode ---
# Find the maximum number of simultaneous rejections at FDR level 5%
r <- closedBY(p, alpha = 0.05)
cat("Largest simultaneous rejection set:", r, "\n")

# The r hypotheses with the smallest p-values form a valid discovery set
discovery_set <- p <= sort(p)[r]
cat("P-values in discovery set:", round(sort(p[discovery_set]), 4), "\n")

# --- Set-checking mode ---
# Check whether a researcher-defined set is a valid simultaneous rejection
candidate_set <- p < 0.01
closedBY(p, set = candidate_set, alpha = 0.05)

# --- Exact vs. approximate ---
r_exact  <- closedBY(p, alpha = 0.05, approximate = FALSE)
r_approx <- closedBY(p, alpha = 0.05, approximate = TRUE)
cat("Exact:", r_exact, "  Approximate:", r_approx, "\n")
```

`closedeBH` *Closed eBH procedure for simultaneous FDR control*

Description

Applies the closed testing version of the e-BH (e-values Benjamini-Hochberg) procedure. The standard eBH procedure controls the false discovery rate (FDR) at level α but only provides a single set of rejections. The closed eBH procedure provides **simultaneous FDR control**: for every set of hypotheses, it determines whether that set can be reported as discoveries while maintaining FDR control at level α, regardless of which other sets were inspected.

Usage

```
closedeBH(e, set = NULL, alpha = 0.05, approximate = FALSE)
```

Arguments

`e`	Numeric vector of e-values, one per hypothesis. E-values must be non-negative; each is interpreted as the evidence against its null hypothesis. The e-values should have expectation at most 1 under the null hypothesis.
`set`	Optional subsetting vector for `e` (logical, index or negative index), indicating which hypotheses belong to the set to be checked for mean consistency. If `NULL` (the default), the function instead returns the size of the largest mean-consistent set.
`alpha`	Numeric scalar in $[0, 1]$. The target FDR level. Defaults to `0.05`.
`approximate`	Logical. If `FALSE` (the default), uses an exact algorithm that is guaranteed to find the largest mean-consistent set. If `TRUE`, uses a faster approximate algorithm that may occasionally return a smaller set. The approximate method is recommended for exploratory analyses or large inputs where computation time is a concern.

Details

The closed eBH procedure is based on the concept of **mean consistency**. A set R of hypotheses is mean consistent — and therefore a valid simultaneous rejection — if and only if:

$$\frac{1}{|S|}\sum_{i\in S} e_i \;\geq\; \frac{|S\cap R|}{|R|\cdot\alpha}$$

is satisfied jointly for all $S \subseteq [m]$, when m hypotheses are tested. In practice, this condition guarantees that R is a valid closed-testing rejection, providing post-hoc FDR control: you may report any mean-consistent set as your discovery set without inflating the FDR above α, even if the choice of set was data-driven.

The function has two modes:

- **Set-checking mode** (when `set` is supplied): Returns `TRUE` if the specified set is mean consistent (i.e., can be reported as a valid simultaneous rejection at level α), and `FALSE` otherwise.
- **Discovery mode** (when `set = NULL`): Returns the size r of the largest mean-consistent set. The r hypotheses with the largest e-values always form one such set. This gives the maximum number of hypotheses that can be reported while maintaining simultaneous FDR control.

Note that mean consistency is not a monotone property: a set of size r being mean consistent does not imply that all smaller sets are as well. The exact algorithm therefore checks all set sizes, while the approximate algorithm (`approximate = TRUE`) uses a faster bisection strategy that may occasionally underestimate the largest consistent set.

Value

- If `set` is supplied: a single logical value. `TRUE` indicates that the specified set is mean consistent and can be reported as a simultaneous rejection at FDR level α. `FALSE` indicates it cannot.

- If `set = NULL`: a single non-negative integer r. The r hypotheses with the largest e-values form a valid simultaneous rejection set. A return value of `0` means no non-empty set can be rejected.

References

Benjamini, Y., & 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.

Wang, R., & Ramdas, A. (2022). False discovery rate control with e-values. *Journal of the Royal Statistical Society: Series B*, 84(3), 822–852.

Xu, Z., Solari, A., Fischer, L., de Heide, R., Ramdas, A., & Goeman, J. (2025). Bringing closure to false discovery rate control: A general principle for multiple testing. arXiv preprint arXiv:2509.02517.

See Also

`closedBH()`, `closedBY()`, `closedBY()`, eClosure

Examples

```
set.seed(42)
# 20 null hypotheses (e ~ Exp(1)) and 10 non-nulls (e ~ Exp(0.1), larger on average)
e <- c(rexp(20, rate = 1), rexp(10, rate = 0.1))

# --- Discovery mode ---
# Find the maximum number of simultaneous rejections at FDR level 5%
r <- closedeBH(e, alpha = 0.05)
cat("Largest simultaneous rejection set:", r, "\n")

# The r hypotheses with the largest e-values form a valid discovery set
discovery_set <- e >= sort(e, decreasing = TRUE)[r]
cat("E-values in discovery set:", round(sort(e[discovery_set], decreasing = TRUE), 2), "\n")

# --- Set-checking mode ---
# Check whether a researcher-defined set is a valid simultaneous rejection
candidate_set <- e > 3
closedeBH(e, set = candidate_set, alpha = 0.05)

# --- Exact vs. approximate ---
r_exact  <- closedeBH(e, alpha = 0.05, approximate = FALSE)
r_approx <- closedeBH(e, alpha = 0.05, approximate = TRUE)
cat("Exact:", r_exact, "  Approximate:", r_approx, "\n")
```

`closedSu` *Closed Su procedure for simultaneous FDR control*

Description

Applies the closed testing improvement of the Su (2018) procedure. The standard Su procedure controls the false discovery rate (FDR) at level α under the PRDN assumption but only provides a single set of rejections. The closed Su procedure provides simultaneous FDR control: for every set of hypotheses, it determines whether that set can be reported as discoveries while maintaining FDR control at level α, regardless of which other sets are inspected.

Usage

```
closedSu(p, set = NULL, alpha = 0.05, approximate = FALSE)
```

Arguments

`p`	Numeric vector of p-values, one per hypothesis.
`set`	Optional subsetting vector for `p` (logical, positive index, or negative index). If `NULL` (the default), returns the size of the largest closed-Su-significant set.
`alpha`	Numeric scalar in $[0, 1]$. FDR level. Default `0.05`.
`approximate`	Logical. If `FALSE` (the default), uses an exact algorithm that checks every set size from largest to smallest and is guaranteed to find the largest closed-Su-significant set. If `TRUE`, uses a faster bisection strategy seeded by the Su lower bound (BH at α/ℓ_α). The approximate method may occasionally underestimate the largest significant set, but is recommended for exploratory analyses or large inputs where computation time is a concern.

Details

A set R of hypotheses is closed-Su-significant — and therefore a valid simultaneous rejection — if and only if, for every $1 \le u \le |R|$ and $0 \le v \le m - |R|$, the combined sorted vector $t_1 \ge \cdots \ge t_{u+v}$ of the u largest p-values in R and the v largest p-values outside R satisfies

$$\min_i \{t_i \cdot (u+v) - c \cdot (u+v-i+1)\} \le 0, \quad c = \frac{|R|\alpha}{u \cdot \ell_\alpha},$$

where $\ell_\alpha = -W_{-1}(-\alpha/e)$ is the Lambert W correction factor (Xu et al., 2025, Section 6.2).

The function has two modes:

- **Set-checking mode** (when `set` is supplied): Returns `TRUE` if the specified set is closed-Su-significant, and `FALSE` otherwise.
- **Discovery mode** (when `set = NULL`): Returns the size r of the largest closed-Su-significant set consisting of the r smallest p-values.

Value

- If `set` is supplied: a single logical (`TRUE`/`FALSE`).
- If `set = NULL`: a single non-negative integer r (0 = no rejection).

References

Su, W. J. (2018). The FDR-Linking Theorem. arXiv:1812.08965.

Xu, Z., Solari, A., Fischer, L., de Heide, R., Ramdas, A., & Goeman, J. (2025). Bringing closure to false discovery rate control. arXiv:2509.02517.

See Also

`closedBH()`, `closedBY()`, `closedeBH()`, eClosure

Examples

```
set.seed(42)
p <- c(runif(20), rbeta(10, 0.1, 1))

# Discovery mode
r <- closedSu(p, alpha = 0.05)

# Set-checking mode
closedSu(p, set = p < 0.01, alpha = 0.05)
```

`eClosure` *eClosure: Multiple testing methods based on the e-Closure principle*

Description

Provides functions for applying the closed Benjamini-Hochberg (BH), closed e-Benjamini-Hochberg (eBH), closed Su and closed Benjamini-Yekutieli (BY) procedures within the e-Closure framework for simultaneous FDR control in multiple hypothesis testing.

Main functions

The two core functions are:

- `closedBH` — applies the closed BH procedure
- `closedeBH` — applies the closed eBH procedure
- `closedBY` — applies the closed BY procedure
- `closedSu` — applies the closed Su procedure

Author(s)

Jelle Goeman

References

Xu, Z., Solari, A., Fischer, L., de Heide, R., Ramdas, A., & Goeman, J. (2025). Bringing closure to false discovery rate control: A general principle for multiple testing. arXiv preprint arXiv:2509.02517.

Goeman, J. (2026). A Uniform Improvement of the Benjamini-Hochberg Procedure via e-Closure. arXiv preprint arXiv:2606.01854.

Index

closedBH, 2, *9*
closedBH(), *5*, *7*, *9*
closedBY, 3, *9*
closedBY(), *3*, *7*, *9*
closedeBH, 5, *9*
closedeBH(), *3*, *5*, *9*
closedSu, 7, *9*
closedSu(), *3*, *5*

eClosure, *5*, *7*, *9*, 9
eClosure-package *(*eClosure*)*, 9

p.adjust(), *3*