Package 'eClosure'

March 9, 2026

Type Package

Title Methods Based on the e-Closure Principle

Version 0.9.2

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 e-Benjamini-Hochberg 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

RoxygenNote 7.3.3

Author Jelle Goeman [aut, cre]

Repository CRAN

Date/Publication 2026-03-09 11:00:09 UTC

Contents

---

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

[p.adjust()](#) for standard p-value-based non-simultaneous multiple testing corrections, including the BY procedure (method = "BY"). [closedBH()](#) for the analogous procedure based on e-values.

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

`p.adjust()` for non-simultaneous multiple testing corrections.

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")
```

---

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

Description

Provides functions for applying the closed e-Benjamini-Hochberg (eBH) 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:

- closedeBH — applies the closed eBH procedure
- closedBY — applies the closed BY 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.

Index