Package 'aucnull'

November 17, 2026

Type Package

Title Heller Test of Equal AUCs for a New Marker in Nested Models

Version 0.2.2

Description Tests whether new markers improve the area under the receiver operating characteristic curve beyond established markers in nested models, using the maximum rank correlation estimator and the weighted chi-square null distribution of Heller, Seshan, Moskowitz and Gonen (2017) <doi:10.1093/biostatistics/kxw045>. The weights are estimated from the bootstrap covariance of the rank correlation estimate (or from a sandwich formula) and the curvature of the kernel smoothed area. The null hypothesis is that the best linear ordering of subjects does not use the new markers, and the test keeps its asymptotic level under any strictly increasing link. Clustered observations are handled by a cluster bootstrap. The likelihood ratio test, Wald tests with a robust covariance and the DeLong test are computed alongside for comparison.

License MIT + file LICENSE

Encoding UTF-8

Depends R (>= 4.1.0)

Imports stats

Suggests testthat (>= 3.0.0), pROC, MASS, survival, mlbench

Config/testthat/edition 3

Config/roxygen2/version 8.0.0

NeedsCompilation no

Author Satoshi Nakamura [aut, cre]

Maintainer Satoshi Nakamura <s-nakamura@med.mie-u.ac.jp>

Repository CRAN

Date/Publication 2026-11-17 19:10:02 UTC

Contents

aucnull_test 2
delong_test 4
mrc_fit 5

print.aucnull_test 5
pweightedchisq 6
smooth_auc 6

Index **8**

`aucnull_test` *Heller test of the null hypothesis of equal AUCs for new markers in nested models*

Description

Tests whether adding the new markers `Z` to the established markers `X` improves the area under the ROC curve of the best linear predictor. The statistic is `2 n (A_full - A_reduced)` with both areas obtained by maximum rank correlation on a kernel smoothed objective. Under the null hypothesis that the best linear ordering does not use `Z`, the statistic converges to a weighted sum of independent chi-square variables with one degree of freedom (Heller et al. 2017, Theorem 1); the weights are the eigenvalues of `-V_gamma %*% solve(D^{gamma gamma})`, where `V_gamma` is the asymptotic covariance of the rank correlation estimate of the coefficients of `Z` and `D` is the Hessian of the smoothed area. Because the weights do not involve the link, the test keeps its asymptotic level for any strictly increasing link, which is not true of the likelihood ratio test when the logistic link is misspecified.

Usage

```
aucnull_test(
  y,
  X,
  Z,
  method = c("boot", "plugin"),
  B = 200,
  expo = 1/3,
  n_sim = 1e+05,
  seed = NULL,
  compare = TRUE,
  cluster = NULL,
  hessian_threshold = -0.01
)
```

Arguments

`y`	binary outcome (0/1 or logical).
`X`	established markers: a numeric vector or matrix (n by p). The first column is the anchor whose coefficient is fixed at one; it must be continuous. Its orientation is chosen by fitting the reduced model with the anchor and with its negative and keeping the orientation with the larger smoothed area (for p = 1 this amounts to an empirical AUC of at least 0.5). Factors must be converted to numeric columns beforehand.

Z	new markers: a numeric vector or matrix (n by q).
method	"boot" or "plugin" for the covariance of the rank correlation estimate.
B	number of bootstrap replicates for method = "boot".
expo	bandwidth exponent for the statistic and the curvature; the bandwidth is sd(X[,1]) * n^(-expo).
n_sim	number of Monte Carlo draws used to evaluate the weighted chi-square distribution when q > 1.
seed	optional seed for the bootstrap; the random number generator state is restored on exit.
compare	logical; if TRUE, the likelihood ratio test of the logistic model, the Wald test with a robust (cluster robust when cluster is given) covariance, the Wald test of Z after a restricted cubic spline in each established marker with the same robust covariance, and the DeLong test on the fitted linear predictors are also returned for comparison. With cluster the likelihood ratio and DeLong tests, which assume independent observations, are not returned.
cluster	optional vector of cluster identifiers of length n.
hessian_threshold	identifiability threshold on the standardized scale: if the smallest eigenvalue of the curvature of the smoothed area in the coefficients of Z after profiling the other coefficients, -solve(solve(D)[gamma, gamma]), is below -hessian_threshold (that is, below 0.01 by default; for one established and one new marker this is the condition that the Hessian exceeds -0.01), the AUC criterion is treated as unidentified and the P value is NA.

Details

Two estimators of V_gamma are offered. "boot" (recommended) uses n times the bootstrap covariance of the rank correlation estimate. "plugin" uses the sandwich form solve(D) Sigma_g solve(D) with Sigma_g estimated from the per observation gradients of the Hajek projection of the smoothed area, computed at the bandwidth sd(X[,1]) n^(-1/3) regardless of expo.

Clustered data (for example several lesions per patient) are handled through cluster: the bootstrap then resamples clusters, the plugin covariance sums the projections within clusters, and the comparison tests use a cluster robust covariance. The limiting distribution is unchanged; the asymptotics are in the number of clusters, so the weights are poorly estimated when there are few clusters. Rows with missing values in y, X, Z or cluster are dropped with a message.

Value

an object of class "aucnull_test": a list with the statistic, the weights lambda, the p value, the rank correlation estimates, the two smoothed areas, the number of clusters, and (if compare = TRUE) the comparison tests.

References

Heller G, Seshan VE, Moskowitz CS, Gonen M (2017). Inference for the difference in the area under the ROC curve derived from nested binary regression models. Biostatistics 18, 260-274.

Han AK (1987). Non-parametric analysis of a generalized regression model. Journal of Econometrics 35, 303-316.

Sherman RP (1993). The limiting distribution of the maximum rank correlation estimator. Econometrica 61, 123-137.

Demler OV, Pencina MJ, D'Agostino RB (2012). Misuse of DeLong test to compare AUCs for nested models. Statistics in Medicine 31, 2577-2587.

Examples

```
set.seed(1)
n <- 200
x <- rnorm(n); z <- rnorm(n)
y <- rbinom(n, 1, plogis(1.2 * x))           # z carries no information
aucnull_test(y, x, z, B = 20)                 # use the default B = 200 in practice

# clustered data: two observations per cluster
id <- rep(1:100, each = 2)
aucnull_test(y, x, z, B = 20, cluster = id)
```

`delong_test` *DeLong test for two correlated AUCs (paired design)*

Description

Provided for comparison only. It is valid for two fixed scores on the same subjects and is not valid for the linear predictors of nested models fitted on the same data (Demler et al. 2012).

Usage

```
delong_test(s1, s2, y)
```

Arguments

`s1`, `s2`	scores.
`y`	binary outcome.

Value

list with the two empirical areas, their difference, the standard error, the z statistic and the two sided p value.

mrc_fit *Maximum rank correlation fit with a smoothed objective*

Description

Maximizes the smoothed area smooth_auc of the linear predictor X[,1] + X[,-1] %*% eta + Z %*% gamma over (eta, gamma). The coefficient of the first established marker is fixed at one for identifiability (Han 1987; Sherman 1993). Pass Z = NULL for the reduced model.

Usage

```
mrc_fit(y, X, Z = NULL, h, start = NULL, reltol = 1e-10)
```

Arguments

y	binary outcome.
X	numeric matrix (n by p) of established markers, p >= 1. The first column should be oriented so that larger values indicate a case.
Z	numeric matrix (n by q) of new markers, or NULL.
h	bandwidth.
start	optional starting values of length p - 1 + q; by default the logistic estimates scaled to the normalization are supplied by aucnull_test, and a grid search is used when none are given.
reltol	convergence tolerance passed to optim. With one free coefficient the maximization is a bounded scalar search (Brent) on an interval of half width two (on the scale of the anchor) around the start, widened while the maximum lies on the boundary; with several free coefficients a quasi Newton (BFGS) search is used.

Value

a list with theta (named coefficient vector, excluding the fixed first coefficient), auc (maximized smoothed area), the objective function f (returns the smoothed area at a parameter vector) and convergence.

print.aucnull_test *Print method for* aucnull_test *objects*

Description

Print method for aucnull_test objects

Usage

```
## S3 method for class 'aucnull_test'
print(x, digits = 4, ...)
```

Arguments

x	an object returned by aucnull_test.
digits	number of significant digits (currently unused; four are shown).
...	further arguments, ignored.

Value

x, invisibly.

---

pweightedchisq *Distribution function of a weighted sum of chi-square(1) variables*

Description

Returns P(sum lambda_j chi2_1j > x). Exact for a single weight; Monte Carlo otherwise.

Usage

```
pweightedchisq(x, lambda, n_sim = 1e+05)
```

Arguments

x	value of the statistic.
lambda	vector of weights.
n_sim	number of Monte Carlo draws when length(lambda) > 1.

Value

upper tail probability.

---

smooth_auc *Kernel smoothed area under the ROC curve*

Description

Computes the kernel smoothed Mann and Whitney statistic $A_h(s) = (n_0 n_1)^{-1} \sum_{i:y_i=1} \sum_{j:y_j=0} \Phi((s_i - s_j)/h)$, as used by Heller et al. (2017) to obtain a differentiable objective.

Usage

```
smooth_auc(score, y, h)
```

Arguments

`score`	numeric vector of scores (higher means more likely to be a case).
`y`	binary outcome (0/1, or logical).
`h`	bandwidth. With `h = 0` the empirical (unsmoothed) area is returned.

Value

a single number in (0, 1).

Index

aucnull_test, 2, *5*, *6*

delong_test, 4

mrc_fit, 5

print.aucnull_test, 5
pweightedchisq, 6

smooth_auc, *5*, 6