

Package: regstat (via r-universe)

September 25, 2026

Title An Exact Test for a Change in Covariance (Dependence) Structure

Version 0.1.0

Description An exact finite-sample test for whether two groups share a covariance matrix, the omnibus form of the differential-network question. Under the Gaussian null the likelihood-ratio statistic has a distribution given by the real Jacobi ensemble that is free of the unknown common covariance, so a single Monte-Carlo calibration at the identity serves every covariance with no estimate of the nuisance covariance; this is the property that survives the dimension barrier, where estimating the covariance is hardest. The max-type high-dimensional test of Cai, Liu and Xia (2013) <doi:10.1080/01621459.2012.758041> is provided for comparison. A pure-C back-end does the numerics and also backs the 'Python' package 'regstat'.

URL <https://github.com/mtloots/regstat>

BugReports <https://github.com/mtloots/regstat/issues>

License GPL-3

Encoding UTF-8

Imports stats

NeedsCompilation yes

Config/roxygen2/version 8.0.0

Config/testthat/edition 3

Suggests knitr, rmarkdown, testthat (>= 3.0.0)

VignetteBuilder knitr

Repository <https://mtloots.r-universe.dev>

Date/Publication 2026-09-14 14:05:09 UTC

RemoteUrl <https://github.com/mtloots/regstat>

RemoteRef HEAD

RemoteSha d50ad55ca4eedcb6cda5e2148129212e2a63ee92

RemoteSubdir regstat

Contents

clx_stat	2
clx_test	2
cov_logdiff	3
cov_M	4
cov_null	4
cov_pdet	5
cov_pexact	6
cov_test	7
diffnet_test	7
Index	9

clx_stat	<i>Cai-Liu-Xia max-type covariance statistic (C back-end)</i>
----------	---

Description

Cai-Liu-Xia max-type covariance statistic (C back-end)

Usage

clx_stat(XA, XB)

Arguments

XA, XB numeric data matrices with the same number of columns.

Value

a single numeric value: the largest standardised squared difference over all $p(p + 1)/2$ covariance entries. It is the statistic of the max-type test, so it responds to a change in a few entries rather than to a diffuse change across many.

clx_test	<i>Max-type two-sample covariance test with the extreme-value calibration</i>
----------	---

Description

The Cai-Liu-Xia (2013) test as used out of the box. Its extreme-value calibration is accurate at large dimension but inflates near the dimension barrier; see [cov_test](#) for the exact, covariance-free alternative.

Usage

clx_test(XA, XB, alpha = 0.05)

Arguments

XA, XB	numeric data matrices with the same number of columns.
alpha	nominal level.

Value

an object of class "htest": a list with statistic (the maximum standardised squared difference), p.value from the extreme-value null, method and data.name. It prints as a standard hypothesis test.

cov_logdiff	<i>Log-domain differential network (C back-end)</i>
-------------	---

Description

Computes the inversion-invariant differential network $D = \log \text{cov}(XB) - \log \text{cov}(XA)$ between two groups. Because the logarithm of a precision matrix is minus the logarithm of the covariance, D is the same object whether dependence is read through covariances or precisions, so it dissolves the covariance-versus-precision choice of the differential network. D is symmetric; its zero pattern is the changed dependence structure.

Usage

```
cov_logdiff(XA, XB)
```

Arguments

XA, XB	numeric data matrices with the same number of columns.
--------	--

Value

a numeric matrix of dimension $p \times p$, symmetric: the log-domain contrast D between the two covariance structures. An entry is zero when the corresponding dependence is unchanged between the groups, so the matrix is read as the estimated differential network itself rather than as a test statistic.

Examples

```
## the log-domain differential network: symmetric, and antisymmetric in its
## arguments, so it reads the same in covariances or precisions
set.seed(1)
XA <- matrix(rnorm(180), 60, 3)
XB <- matrix(rnorm(150), 50, 3) %*% diag(c(1, 1.6, 0.7))
D <- cov_logdiff(XA, XB)
max(abs(D + cov_logdiff(XB, XA)))
```

cov_M	<i>Box's M statistic for equality of two covariance matrices (C back-end)</i>
-------	---

Description

Box's M statistic for equality of two covariance matrices (C back-end)

Usage

```
cov_M(XA, XB)
```

Arguments

XA	numeric data matrix, group A (rows observations, columns variables).
XB	numeric data matrix, group B, with the same number of columns.

Value

a single numeric value: the likelihood-ratio statistic M for equality of the two covariance matrices, on the scale for which the null law is the real Jacobi ensemble. Larger values are stronger evidence against equal covariance. NaN if either group scatter is singular, which happens when a group has fewer observations than dimensions.

Examples

```
## Box's M vanishes when nothing differs, and is invariant to a common basis change
set.seed(1)
XA <- matrix(rnorm(180), 60, 3)
XB <- matrix(rnorm(150), 50, 3) %*% diag(c(1, 1.6, 0.7))
c(same = cov_M(XA, XA), different = cov_M(XA, XB))
```

cov_null	<i>Draws from the covariance-free null of M (C back-end)</i>
----------	--

Description

Simulates the exact Gaussian null of `cov_M` at the given dimensions. Because the null is free of the common covariance, these draws (made at the identity) are the reference law at every covariance.

Usage

```
cov_null(p, nA, nB, B = 8000, seed = 1L)
```

Arguments

p	number of variables.
nA, nB	per-group sample sizes.
B	number of draws.
seed	integer seed for the self-contained generator.

Value

a numeric vector of length B: independent draws of the statistic M under the null of equal covariance. Because that null is free of the unknown common covariance, the draws may be taken at the identity and reused for any covariance; the vector is the Monte Carlo reference distribution against which an observed M is compared.

cov_pdet	<i>Exact null of the determinant functional, and exact power against proportional alternatives</i>
----------	--

Description

Upper-tail probability of $D = \log(\det(A2)/\det(A1))$ under H_0 , for independent Wisharts with nu1 and nu2 degrees of freedom in p dimensions; the null is covariance-free, and its characteristic function is a pure product of gamma ratios by the Bartlett decomposition, inverted by the same safeguarded Gil-Pelaez quadrature as cov_pexact. Under the proportional alternative $\Sigma_2 = \lambda * \Sigma_1$ the statistic shifts exactly by $p * \log(\lambda)$, so cov_powdet returns exact, deterministic power for the determinant chart at threshold crit.

Usage

```
cov_pdet(x, nu1, nu2, p)

cov_powdet(lambda, crit, nu1, nu2, p)
```

Arguments

x	threshold for the upper-tail probability.
nu1, nu2	within-group degrees of freedom (sample size minus one).
p	number of variables.
lambda	scale factor of the proportional alternative.
crit	chart threshold on the D scale.

Value

a single numeric value in $[0, 1]$ from each function. `cov_pdet` returns the exact upper-tail probability $P(D > x)$ of the determinant statistic under the null of equal covariance, so it is a p-value. `cov_powdet` returns the exact power of the determinant chart at threshold `crit` against the proportional alternative $\Sigma_2 = \lambda \Sigma_1$, evaluated at each element of `lambda`, so its length is that of `lambda`. The power is deterministic, not simulated, because the alternative shifts the statistic by exactly $p \log \lambda$.

<code>cov_pexact</code>	<i>Exact right-tail probability of Box's M (C back-end)</i>
-------------------------	--

Description

Evaluates $P(M > m)$ under the covariance-free null deterministically, by numerical inversion of the exact characteristic function – no simulation. Because the null is covariance-free the result depends only on the dimensions.

Usage

```
cov_pexact(m, nu1, nu2, p)
```

Arguments

<code>m</code>	observed statistic value.
<code>nu1, nu2</code>	within-group degrees of freedom (sample size minus one).
<code>p</code>	number of variables.

Value

a single numeric value in $[0, 1]$: the exact upper-tail probability $P(M > m)$ under the null of equal covariance, obtained by inverting the characteristic function rather than by simulation. Used directly as a p-value.

Examples

```
## the null tail by inversion of the exact characteristic function
vapply(c(5, 15, 35), function(m) cov_pexact(m, 59, 49, 3), 0)
```

cov_test

*Exact test for a change in covariance structure***Description**

Tests H_0 that two groups share a covariance matrix, using the likelihood-ratio statistic M and its covariance-free null. By default the null is evaluated exactly and deterministically by inverting the characteristic function (`method = "exact"`); `method = "mc"` uses a Monte Carlo null instead. Either way the test needs no estimate of the common covariance, its advantage over high-dimensional competitors near the dimension barrier where that covariance is hardest to estimate.

Usage

```
cov_test(XA, XB, method = c("exact", "mc"), B = 8000, seed = 1L)
```

Arguments

XA, XB	numeric data matrices with the same number of columns.
method	"exact" (CF inversion) or "mc" (Monte Carlo).
B	number of null draws when <code>method = "mc"</code> .
seed	integer seed when <code>method = "mc"</code> .

Value

an object of class "htest": a list with `statistic` (the observed M), `p.value`, `method` naming which null was used, and `data.name`. It prints as a standard hypothesis test. The p-value is exact when `method = "exact"` and Monte Carlo otherwise.

Examples

```
set.seed(1)
XA <- matrix(rnorm(180), 60, 3)
XB <- matrix(rnorm(150), 50, 3) %*% diag(c(1, 1.6, 0.7))
cov_test(XA, XB, method = "exact")$p.value
```

diffnet_test

*Two-sample test for a change in dependence structure via the log-domain contrast***Description**

Tests H_0 that two groups share a covariance (equivalently precision) structure, using the inversion-invariant contrast `cov_logdiff` and its bootstrap covariance. The statistic is the quadratic form in the vectorised upper triangle of D , referred to a chi-squared law on $p(p+1)/2$ degrees of freedom. Unlike the covariance and precision differential networks, the estimand tested is unambiguous.

Usage

```
diffnet_test(XA, XB, B = 400, seed = 1L)
```

Arguments

XA, XB	numeric data matrices with the same number of columns.
B	number of bootstrap resamples for the covariance of D.
seed	integer seed for the resampling.

Value

an object of class "htest": a list with statistic (the quadratic form T), parameter (its degrees of freedom), p.value from the chi-squared null, method and data.name. It prints as a standard hypothesis test. A small p-value says the two dependence structures differ somewhere, without saying where.

Index

clx_stat, [2](#)
clx_test, [2](#)
cov_logdiff, [3](#), [7](#)
cov_M, [4](#), [4](#)
cov_null, [4](#)
cov_pdet, [5](#)
cov_pexact, [6](#)
cov_powdet (cov_pdet), [5](#)
cov_test, [2](#), [7](#)

diffnet_test, [7](#)