

Package: matchednull (via r-universe)

July 23, 2026

Title Matched-Null Tests for Cluster-Count Claims

Version 0.1.0.9000

Description Builds matched nulls for cluster-count claims: synthetic twins of a dataset that preserve every marginal distribution and the full correlation matrix while containing no cluster structure by construction. A reported number of clusters or ``types" can then be tested against what the data's own margins and covariance already produce, using any clustering pipeline. A t-copula option adds tail dependence to the null, so that an apparent excess of clusters can be checked against a heavier-tailed alternative before it is read as evidence of types. Implements the matched-null procedure of Meng (2026) ``Types Without Taxa" <<https://osf.io/2ekcg>>.

License MIT + file LICENSE

Encoding UTF-8

Roxygen list(markdown = TRUE)

RoxygenNote 7.3.3

Imports stats

Suggests mclust, testthat (>= 3.0.0), knitr, rmarkdown

VignetteBuilder knitr

Config/testthat/edition 3

URL <https://github.com/haomeng797-ship-it/matchednull>

BugReports <https://github.com/haomeng797-ship-it/matchednull/issues>

Repository <https://haomeng797-ship-it.r-universe.dev>

Date/Publication 2026-07-23 00:12:17 UTC

RemoteUrl <https://github.com/haomeng797-ship-it/matchednull>

RemoteRef HEAD

RemoteSha 7008fbead13d7eeb60493fb31ddadae95fc70eaf

Contents

copula_null	2
matched_null_test	3
Index	5

copula_null	<i>Draw one matched null ("null twin") of a dataset</i>
-------------	---

Description

Builds a synthetic twin of `x` that preserves every marginal distribution exactly and the correlation matrix to within sampling error, while containing no cluster structure by construction. With `copula = "gaussian"` (the default) all dependence in the twin is Gaussian. Clustering found in real data but not in its twins must come from structure beyond the margins and covariance; clustering found equally in both was never more than the data's shape.

Usage

```
copula_null(x, copula = c("gaussian", "t"), df = 8, ridge = 1e-06)
```

Arguments

<code>x</code>	A numeric matrix or data frame (rows = observations, columns = variables), complete cases only, at least two rows and two columns.
<code>copula</code>	Dependence family of the twin: "gaussian" (default) or "t".
<code>df</code>	Degrees of freedom of the t copula, a single positive number (default 8; used only when <code>copula = "t"</code>). Smaller is heavier-tailed; <code>df = 3</code> is a hard stress test, <code>df = 8</code> a moderate one.
<code>ridge</code>	Small value added to the diagonal of the correlation matrix only if it is not positive definite (default 1e-6).

Details

The construction is a rank reordering in the tradition of Iman and Conover (1982): draw a Gaussian sample with the data's correlation matrix, then replace each column with the sorted real values laid down in the rank order of the Gaussian column. Every real value is reused exactly once per column, which is why the margins match exactly; the rank correlation is matched exactly and the Pearson correlation to within sampling error.

With `copula = "t"` the same construction is driven by a multivariate t sample instead: margins and correlations are preserved as before, but the twin also carries tail dependence, so extreme values across variables arrive together, the more strongly the smaller `df`. A t twin is still a single population with no clusters. Its use is as a stress test: a verdict of "exceeds the Gaussian null" that a t twin reproduces was heavy-tailed dependence, not types.

Value

A numeric matrix of the same dimensions as `x`: one matched-null draw. Each column contains exactly the values of the corresponding column of `x`, rearranged.

References

Iman, R. L., & Conover, W. J. (1982). A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics - Simulation and Computation*, 11(3), 311-334.

Examples

```
set.seed(1)
x <- matrix(rnorm(200 * 3), 200, 3) %%% chol(matrix(c(1, .5, .3,
                                                    .5, 1, .4,
                                                    .3, .4, 1), 3, 3))

twin <- copula_null(x)
# margins identical:
all(sort(twin[, 1]) == sort(x[, 1]))
# correlations close:
round(cor(x) - cor(twin), 2)

# a heavier-tailed twin for stress-testing:
stress <- copula_null(x, copula = "t", df = 3)
all(sort(stress[, 1]) == sort(x[, 1]))
```

matched_null_test	<i>Test a cluster count against matched nulls</i>
-------------------	---

Description

Runs the user's own clustering pipeline on the real data and on `R` matched-null twins of it (see [copula_null\(\)](#)), and asks whether the real result exceeds what the twins, which share the data's margins and covariance but contain no cluster structure, produce. The pipeline is supplied as a function, so any procedure that returns a number, a BIC-selected mixture, a k-means heuristic, a published typology's own workflow, can be tested unchanged.

Usage

```
matched_null_test(
  x,
  cluster_fn,
  R = 200,
  probs = c(0.025, 0.975),
  copula = c("gaussian", "t"),
  df = 8,
  ridge = 1e-06
)
```

Arguments

<code>x</code>	A numeric matrix or data frame, complete cases only.
<code>cluster_fn</code>	A function taking a data matrix and returning a single number: the statistic to be tested, typically the selected number of clusters (but any scalar summary of clustering strength works).
<code>R</code>	Number of matched-null twins to draw (default 200).
<code>probs</code>	Lower and upper quantiles of the null distribution used for the interval verdict (default <code>c(.025, .975)</code>).
<code>copula, df</code>	Dependence family of the twins and <code>t</code> degrees of freedom, passed to <code>copula_null()</code> . Defaults to "gaussian".
<code>ridge</code>	Passed to <code>copula_null()</code> .

Details

A verdict of "exceeds the null" under the default Gaussian twins licenses only "structure beyond margins and correlations", it does not by itself license types: heavy-tailed dependence also exceeds a Gaussian null. To separate the two, rerun with `copula = "t"` (df 8, then 3). A result that survives the `t` twins as well is harder to attribute to tails; a result the `t` twins reproduce was tail dependence, not types.

Value

An object of class "matched_null_test": a list with the real statistic (`real`), the null draws (`null`), the null interval (`interval`), a one-sided Monte Carlo p-value for exceedance (`p_exceed`), the interval verdict (`within`), and the null family used (`copula`, with `df` when it is "t"). Reproducibility is the caller's: set a seed before calling.

Examples

```
if (requireNamespace("mclust", quietly = TRUE)) {
  suppressPackageStartupMessages(library(mclust))
  # a pipeline: how many components does BIC select?
  pick_k <- function(d) Mclust(d, G = 1:5, verbose = FALSE)$G
  set.seed(42)
  x <- matrix(rnorm(500 * 4), 500, 4) # typeless data
  matched_null_test(x, pick_k, R = 30)

  # stress test of an exceedance verdict against heavier tails:
  # matched_null_test(x, pick_k, R = 30, copula = "t", df = 8)
}
```

Index

`copula_null`, [2](#)

`copula_null()`, [3](#), [4](#)

`matched_null_test`, [3](#)