

Package: renyi (via r-universe)

September 20, 2026

Type Package

Title Renyi Outlier Test

Version 1.0.0

Author Ryan Christ <ryan.christ@yale.edu>

Maintainer Ryan Christ <ryan.christ@yale.edu>

Description renyi implements the Renyi Outlier Test
<[doi:10.48550/arXiv.2411.13542](https://doi.org/10.48550/arXiv.2411.13542)>, an outlier test designed for modern large
scale testing applications, especially where prior information
available. The test combines a vector of independent uniform
p-values into one p-value with power against alternatives where
a small number of p-values are non-null. The test can leverage
prior probabilities/weights specifying which variables are
likely to be outliers and prior estimates of effect size. The
procedure is fast even when the number of initial p-values is
large (e.g. in the millions) and numerically stable even for
very small p-values (e.g. 10^{-300}).

License Apache License (≥ 2)

BugReports <https://github.com/ryanchrist/renyi/issues>

URL <https://ryanchrist.r-universe.dev/renyi>,
<https://github.com/ryanchrist/renyi>

Encoding UTF-8

LazyData true

RoxygenNote 7.3.3

Depends R ($\geq 3.5.0$)

Imports stats, utils

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

Date/Publication 2025-12-15 18:16:09 UTC

RemoteUrl <https://github.com/ryanchrist/renyi>

RemoteRef HEAD

RemoteSha 9390dc6b6c5d2a2a76845c76ed7baa62014969de

Contents

generalized_renyi_transform	2
renyi	3
uniformity_tests	4
Index	7

generalized_renyi_transform
<i>Generalized Renyi Transform</i>

Description

A Generalization of Aldous Renyi’s representation of exponential order statistics

Usage

```
generalized_renyi_transform(x, eta = NULL, zeta = NULL)
```

Arguments

x	a vector of independent exponential random variables of the form $X_j = \eta_j Y_j + \zeta_j$ where each X_j is an independent exponential random variable with rate 1
eta	vector of scale parameters implicit in the construction of x : eta [j] = η_j
zeta	vector of shift parameters implicit in the construction of x : zeta [j] = ζ_j

Details

Maps a vector of shifted and scaled independent exponential random variables to a sequence of standard independent exponential random variables based on the gaps (jumps) between the initial random variables

Value

- a list containing two elements
- ‘exps’** a vector of independent standard exponentials where **exps**[1] is the exponential jump corresponding to **min(x)** and **tail(exps,1)** is the exponential jump corresponding to **max(x)**.
- ‘order’** **order(x)**.

References

Christ, R., Hall, I. and Steinsaltz, D. (2024) "The Renyi Outlier Test", [arXiv:2411.13542](#) . Available at: [doi:10.48550/arXiv.2411.13542](#).

Examples

```
# example code

a <- rchisq(10,1)
b <- rnorm(10)
xx <- a*rexp(10)+b
generalized_renyi_transform(xx, a, b)
```

renyi

Renyi Outlier Test

Description

A fast, numerically precise outlier test for a vector of exact p-values allowing for prior information

Usage

```
renyi(u, k = ceiling(0.01 * length(u)), pi = NULL, eta = NULL)
```

Arguments

u	a vector of p-values
k	a rough upper bound on the number of outliers expected to be present in u
pi	optional vector such that <code>pi[j]</code> is proportional to the probability that <code>u[j]</code> is an outlier. The default, <code>NULL</code> , corresponds to <code>pi = rep_len(1, length(u))</code> .
eta	optional vector proportional to how far outlying we expect <code>u[j]</code> to be given <code>u[j]</code> is an outlier. More precisely, in the common context where each element of <code>u</code> can be thought of as a p-value for testing whether some coefficient β in a linear regression model is zero, we assume <code>eta[j]</code> is proportional to $\mathbb{E}[\beta_j^2 \beta_j \neq 0]$. The default, <code>NULL</code> , corresponds to <code>eta = rep_len(1, length(u))</code> .

Details

The about which p-values are outlying and "how much" of an outlier they are expected to be

Value

a list containing three elements

‘**p_value**’ the p-value returned by the Renyi Outlier Test;

‘**max_k**’ a power of 2 in $2^{(0:k)}$ denoting the number of tail p-values that yielded the most significant signal when running the Renyi Outlier Test;

‘**p_value_k1**‘ the p-value that would be returned by the Renyi Outlier Test assuming $k=1$;

‘**exit_status**‘ a character string describing any problems that may have been encountered during evaluation, ”default is no problems”;

‘**u**‘ the vector of p-values used by the outlier test after adjusting the **u** provided for **pi** and **eta**.

References

Christ, R., Hall, I. and Steinsaltz, D. (2024) ”The Renyi Outlier Test”, [arXiv:2411.13542](#) . Available at: [doi:10.48550/arXiv.2411.13542](#).

Examples

```
# example code

p <- 1e4
u <- runif(p)
u[c(53,88,32)] <- 1e-6 # add a few outliers
renyi(u)$p_value # test for outliers without any prior knowledge
renyi(u,pi=c(rep(1,100),rep(10^-3,p-100)))$p_value # test for outliers with prior knowledge
```

uniformity_tests	<i>Test for Uniformity on $[0,1]$ Using Multiple Statistical Tests</i>
------------------	---

Description

A wrapper function that performs multiple statistical tests to assess whether a numeric vector represents independent draws from a uniform distribution on the interval $[0,1]$. The function combines several complementary approaches including tests based on the Rényi Outlier Test (see [renyi](#)), distribution fitting (Kolmogorov-Smirnov), location (t-test), and normality after transformation (Shapiro-Wilk).

Usage

```
uniformity_tests(u, k = 32)
```

Arguments

u A numeric vector of values assumed to be on $[0,1]$. Each element should represent an independent draw that is being tested for uniformity.

... optional arguments to be passed to the Rényi Outlier Test, see [renyi](#).

Details

The function applies four different statistical tests:

- **Kolmogorov-Smirnov:** Compares the empirical distribution of u to the uniform distribution on $[0,1]$
- **Rényi Outlier Test:** Tests whether there are outlying small entries of u , see [renyi](#).
- **t-test:** Transforms u using the inverse normal CDF and tests whether the mean equals 0 (expected value for standard normal)
- **Shapiro-Wilk:** Tests whether the normal quantile transform $\Phi^{-1}(u)$ follows a standard normal distribution

For the Shapiro-Wilk test, if the sample size exceeds 5,000, the function automatically subsamples to 5,000 quantiles to meet the test's sample size limitations.

All tests return p-values where small values (typically < 0.05) suggest evidence against the null hypothesis of uniformity.

Value

A named list containing p-values from four uniformity tests:

shapiro P-value from Shapiro-Wilk test applied to normal quantile transformed data. Tests whether $\Phi^{-1}(u)$ follows standard normal distribution.

renyi P-value from the Rényi Outlier Test.

t P-value from one-sample t-test testing if mean of $\Phi^{-1}(u)$ equals 0.

ks P-value from Kolmogorov-Smirnov test comparing the empirical distribution to the uniform distribution.

See Also

[ks.test](#), [shapiro.test](#), [t.test](#), [renyi](#)

Examples

```
# Test truly uniform data
uniform_data <- runif(1000)
results1 <- uniformity_tests(uniform_data)
print(results1) # Should show large p-values

# Test non-uniform data (beta distribution)
beta_data <- rbeta(1000, 2, 5)
results2 <- uniformity_tests(beta_data)
print(results2) # Should show small p-values

# Test a data with small u outliers
outlier_data <- c(uniform_data, 1e-5, 5e-6, 1e-6)
results3 <- uniformity_tests(outlier_data)
print(results3) # Should show small p-values

# Test while passing a different argument k to the Rényi Outlier Test
```

```
results4 <- uniformity_tests(outlier_data, k = 4)
print(results4)
```

Index

`generalized_renyi_transform`, [2](#)

`ks.test`, [5](#)

`renyi`, [3](#), [4](#), [5](#)

`shapiro.test`, [5](#)

`t.test`, [5](#)

`uniformity_tests`, [4](#)