

Package: multChernoff (via r-universe)

September 20, 2026

Title Finite-Sample Tail Bound of Likelihood Ratio Test under Multinomial Sampling

Version 1.0.0

Description Computes a finite-sample tail bound for the log-likelihood ratio test (LRT) statistic under multinomial sampling. The resulting bound is used to compute finite-sample conservative p-values and critical values when the standard chi-squared asymptotics can be unreliable. The package also supports multiple independent multinomial trials.

License MIT + file LICENSE

Encoding UTF-8

Roxygen list(markdown = TRUE)

RoxygenNote 7.3.3

URL <https://github.com/richardkwo/multChernoff>

BugReports <https://github.com/richardkwo/multChernoff/issues>

Imports plyr

Suggests knitr, rmarkdown, testthat (>= 3.0.0)

VignetteBuilder knitr

Config/testthat/edition 3

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

Date/Publication 2026-03-24 22:53:15 UTC

RemoteUrl <https://github.com/richardkwo/multchernoff>

RemoteRef HEAD

RemoteSha 3d0117d7e6249fd2169f57f26a059c0db1d02c85

Contents

criticalValue	2
mgfBound	3
tailProbBound	4

criticalValue	<i>Critical value x such that $P(2LRT > x) \leq p$</i>
---------------	---

Description

The LRT is the log-likelihood ratio test statistic, which can be written as

$$LRT = n KL(\hat{p}||p).$$

By the Wilks' theorem, for a fixed k -dimensional probability vector, it holds that

$$2 \times LRT \rightarrow_d \chi_{k-1}^2.$$

This function returns a finite-sample counterpart to `qchisq(p, k-1, lower.tail=FALSE)`. The LRT is also extended to multiple independent multinomial trials. For example, for a (k_1, n_1) -trial and a (k_2, n_2) -trial, we have

$$LRT = n_1 KL(\hat{p}_1||p_1) + n_2 KL(\hat{p}_2||p_2).$$

Usage

```
criticalValue(k, n, p = 0.05, verbose = FALSE)
```

Arguments

<code>k</code>	number of categories (a vector for independent multinomial draws)
<code>n</code>	sample size (a vector for independent multinomial draws)
<code>p</code>	significance level (e.g., 0.05)
<code>verbose</code>	draw the minimizer if TRUE

Value

A finite-sample critical value x such that the bound on $P(2 \times LRT > x)$ is at most p .

Note

For independent multinomial samples, k and n must be of the same length.

See Also

[tailProbBound](#), [mgfBound](#)

Examples

```

n <- 1:40
crit <- sapply(n, function(.n) criticalValue(20, .n, p=0.01))
plot(n, crit)
# chi-squared asymptotic by Wilks' theorem
abline(h=qchisq(0.01, df=20-1, lower.tail = FALSE))
criticalValue(10, 40, p=0.05)
# two independent multinomial trials (k=3, n=4) and (k=12, n=20)
criticalValue(c(3, 4), c(12, 20), p=0.05)

```

mgfBound

*An upper bound on the moment generating function of LRT***Description**

The LRT is the log-likelihood ratio test statistic, which can be written as

$$LRT = n KL(\hat{p}||p),$$

namely the Kullback-Leibler divergence from the empirical probabilities to the true probabilities multiplied by the sample size. $G(\lambda; k, n)$ is a polynomial in λ such that

$$MGF(\lambda; LRT) := E_p[\exp(\lambda \times LRT)] \leq G(\lambda; k, n)$$

holds for every $\lambda \in [0, 1]$ and every p .

Usage

```
mgfBound(k, n, lambda)
```

Arguments

k	number of categories
n	sample size
lambda	number between 0 and 1

Value

A numeric upper bound on the MGF of LRT evaluated at lambda.

See Also

[tailProbBound](#), [criticalValue](#)

Examples

```

mgfBound(k = 5, n = 20, lambda = 0.5)
mgfBound(k = 5, n = 20, lambda = 0) # always 1 at lambda = 0

```

tailProbBound	<i>Tail bound on $P(2 \text{ LRT} > x)$.</i>
---------------	--

Description

The LRT is the log-likelihood ratio test statistic, which can be written as

$$LRT = n KL(\hat{p}||p).$$

By the Wilks' theorem, for a fixed k-dimensional probability vector, it holds that

$$2 \times LRT \rightarrow_d \chi_{k-1}^2.$$

This function returns a finite-sample counterpart to `pchisq(x, k-1, lower.tail=FALSE)`. The LRT is also extended to multiple independent multinomial trials. For example, for a (k_1, n_1) -trial and a (k_2, n_2) -trial, we have

$$LRT = n_1 KL(\hat{p}_1||p_1) + n_2 KL(\hat{p}_2||p_2).$$

Usage

```
tailProbBound(x, k, n, verbose = FALSE)
```

Arguments

x	the value of $2 \times LRT$
k	number of categories (a vector for independent multinomial draws)
n	sample size (a vector for independent multinomial draws)
verbose	draw the minimizer if TRUE

Value

An upper bound on $P(2 \text{ LRT} > x)$, which can be used as a conservative p-value.

Note

For independent multinomial samples, k and n must be of the same length.

See Also

[criticalValue](#), [mgfBound](#)

Examples

```
tailProbBound(20, 7, 50)
pchisq(20, 6, lower.tail=FALSE) # compare with the standard chi-square asymptotic
# two independent multinomial trials (k=3, n=4) and (k=12, n=20)
tailProbBound(12, c(3, 4), c(12, 20))
pchisq(12, 5, lower.tail=FALSE) # compare with the standard chi-square asymptotic
```

Index

`criticalValue`, [2](#), [3](#), [4](#)

`mgfBound`, [2](#), [3](#), [4](#)

`tailProbBound`, [2](#), [3](#), [4](#)