The Journal of Open Source Software

robnptests – An R package for robust two-sample location and dispersion tests

Sermad Abbas [1], **Barbara Brune** [2], **and Roland Fried** [1]

1 TU Dortmund University, Faculty of Statistics, 44221 Dortmund, Germany **2** Technical University of Vienna, Institute of Statistics and Mathematical Methods in Economics, 1040 Vienna, Austria

DOI: 10.21105/joss.04947

Software

- Review
- Repository
- Archive

Editor: Mehmet Hakan Satman

Reviewers:

- @mingzehuang
- @msalibian

Submitted: 14 October 2022
Published: 15 February 2023

Summary

The R (R Core Team, 2022) package `robnptests` is a compilation of two-sample tests selected by two criteria: The tests are (i) robust against outliers and (ii) (approximately) distribution free. Criterion (ii) means that the implemented tests keep an intended significance level and provide a reasonably high power under a variety of continuous distributions. Robustness is achieved by using test statistics that are based on robust location and scale measures.

In what follows, we give a brief description of the package's contents. More details can be found in the introductory vignette of the package, which can be opened by calling `vignette("robnptests")` in the R console, and in the cited papers.

Data situation

We consider two samples of independent and identically distributed (i.i.d.) random variables $X_1, ..., X_m$ and $Y_1, ..., Y_n$, respectively. The underlying distributions are assumed to be continuous with cumulative distribution functions F_X and F_Y.

The tests can be used for either of the following scenarios:

- Two-sample location problem: Assuming that both distributions are equal except that F_Y may be a shifted version of F_X, i.e. $F_X(x) = F_Y(x + \Delta)$ for all $x \in \mathbb{R}$ and some $\Delta \in \mathbb{R}$, or equivalently $X \overset{d}{=} Y - \Delta$, the tests can be used to detect such a shift.
- Two-sample scale problem: In case of a difference only in scale, i.e. $F_X(x) = F_Y(x/\theta)$ for some $\theta > 0$, or equivalently $X \overset{d}{=} Y \cdot \theta$, a transformation of the observations enables to identify differing scale parameters. For more information see `vignette("robnptests")`.

Statement of need

A popular test for the location setting is the two-sample t-test. It is considered to be robust against deviations from the normality assumption because it keeps the specified significance level asymptotically due to the central limit theorem in case of finite variances. However, non-normality can result in a loss of power (Wilcox, 2003). In addition, the t-test is vulnerable to outliers (Fried & Dehling, 2011). Distribution-free tests, like the two-sample Wilcoxon rank-sum test, can be nearly as powerful as the t-test under normality and may have higher power under non-normality. Still, they also can be vulnerable to outliers, particularly for small samples (Fried & Gather, 2007). The two-sample tests in `robnptests` are (approximately) distribution free and, at the same time, robust against outliers. Thus, they are well suited for outlier-corrupted samples from unknown data-generating distributions. At the same time, such tests can be nearly as powerful as popular procedures like the aforementioned t-test or the Wilcoxon test on uncontaminated samples for a somewhat longer computation time.

Figure 1 compares the power of the t-test, the Wilcoxon test and two robust tests. The HL1-test is based on the one-sample Hodges-Lehmann estimator (Hodges & Lehmann, 1963) and the Huber M-test uses Huber's M-estimator (Huber, 1964). We consider a fixed location difference between the samples and a single outlier of increasing size. The power of the t-test decreases to zero, while the loss in power of the Wilcoxon test and both robust tests is small. The robust tests provide a somewhat higher power than the Wilcoxon test and this advantage can become larger when more outliers are involved (Fried & Dehling, 2011).

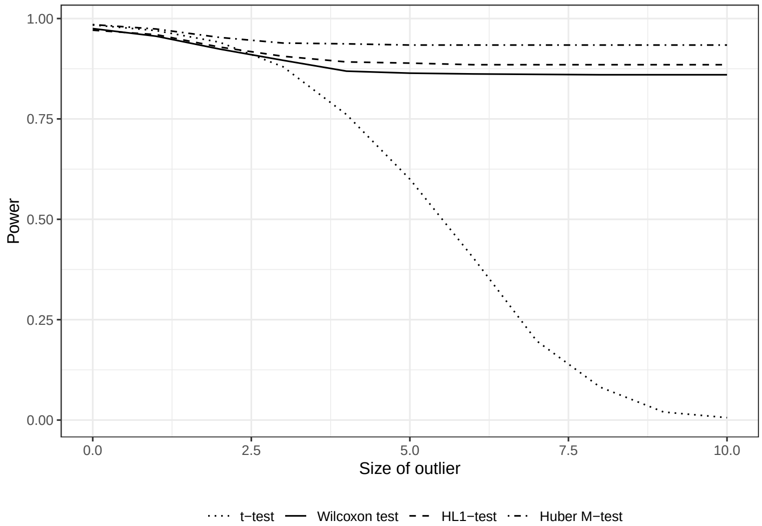

Figure 1: Power of the two-sample t-test, the Wilcoxon rank-sum test, and two robust tests - one based on the one-sample Hodges-Lehmann estimator and one based on Huber's M-estimator - on two samples of size $m = n = 10$ from two normal distributions with unit variance, a location difference of $\Delta = 2$, and an additive single outlier of increasing size.

Common parametric and nonparametric tests for scale differences have similar problems as described above for the location tests. In addition, some nonparametric tests for the scale problem do not cope well with asymmetry. The package `robnptests` uses the idea of applying the robust location tests to transformed observations as proposed by Fried (2012). Such tests are also robust and can obtain good results in terms of power and size under both asymmetry and outlier corruption. However, these tests may be less powerful under symmetry than classical procedures like the Mood test.

Other packages with robust two-sample tests

The package `WRS2` (Mair & Wilcox, 2020) contains a collection of robust two-sample location tests for the heteroscedastic setting. In `robnptests`, we assume homoscedasticity for the location tests. This is because estimating the within-sample dispersion for both samples separately may be unreliable when the sample sizes are small. Equal sample sizes $m = n$ can protect against a deteriorating performance in terms of size and power if we are actually in the heteroscedastic setting (Staudte & Sheather, 1990, p. 179).

The package `nptest` (Helwig, 2021) contains nonparametric versions of the two-sample t-test, realized by using the permutation and randomization principles, as described in the next section, on the t-statistic. This approach has also been studied in Abbas & Fried (2017), and, while being distribution free, the test statistic lacks robustness against outliers.

Implemented two-sample tests

The tests for a location difference are simple ratios inspired by the test statistic of the two-sample t-test. The numerator is a robust estimator for the location difference between the two populations and the denominator is a robust measure for the dispersion within the samples.

The p-value can be computed by using the permutation principle, the randomization principle, or a normal approximation. With the permutation principle, the tests hold the desired significance level exactly at the cost of large computing times even for quite small samples such as $m = n = 10$. The time can be reduced by using a randomization distribution and, even more, by taking advantage of the asymptotic normality of the location-difference estimators. The latter approach, however, is only advisable for large sample sizes $m, n > 30$.

The tests based on the following estimators for the location difference are described in Fried & Dehling (2011):

- The *difference of the sample medians* leads to highly robust tests. However, they are not very powerful under normality due to the low efficiency of the median.
- To obtain more powerful tests under normality, one can use the difference between the *one-sample Hodges-Lehmann estimators* (Hodges & Lehmann, 1963). This may result in less robust tests due to the lower breakdown point.
- The *two-sample Hodges-Lehmann estimator* (Hodges & Lehmann, 1963) leads to robust tests with a higher power under normality than the tests based on the sample median and can achieve similar robustness.

For scaling, we use different estimators based on medians and pairwise differences, see Fried & Dehling (2011) for a detailed description.

In addition, we implemented tests based on M-estimators. This approach to robust location estimation allows for flexibility in how outliers are treated through the specification of the tuning constants of the corresponding ρ-function. We focus on Huber's ρ-function, the bisquare function and Hampel's ρ-function. The measure for the dispersion within the samples is a pooled statistic derived from the asymptotic normality of the M-estimators (Maronna et al., 2019, p. 37ff). Moreover, the package contains Yuen's t-test which uses the difference of *trimmed means* to estimate the location difference and a scale estimator based on the pooled winsorized variances (Yuen & Dixon, 1973).

In case of data with many ties (e.g. caused by discrete sampling), the ties may carry over to the permutation distribution. This can happen in real-world applications when the measurements are rounded or stem from discrete distributions and may lead to a loss in power or conservative tests. Additionally, the robust scale estimators may become zero, so that the test statistic cannot be calculated. Both issues can be addressed by adding random noise from a uniform distribution with a small variance to each observation ("wobbling", see Fried & Gather (2007)).

The following code snippet shows how the tests can be applied to a data set.

```
library(robnptests)

# Generate samples
set.seed(108)
x <- rnorm(10)
y <- rnorm(10)

# Use test based on one-sample Hodges-Lehmann estimator
hl1_test(x = x, y = y, alternative = "two.sided", delta = 0,
         method = "permutation")
#>
#>  Exact permutation test based on HL1-estimator
#>
```

```
#> data:  x and y
#> D = 0.55524, p-value = 0.27
#> alternative hypothesis: true location shift is not equal to 0
#> sample estimates:
#>   HL1 of x   HL1 of y
#>  0.2384959 -0.1140219
```

Here, we use a test based on the one-sample Hodges-Lehmann estimator. By setting `alternative = "two.sided"` and `delta = 0`, we test the null hypothesis $H_0 : \Delta = 0$, i.e. there is no location difference between the populations. In the example above, we use `method = "permutation"` so that the p-value is computed with the permutation principle.

In general, the functions start with the name of the underlying location-difference estimator and have several arguments to customize the test.

More examples on how to use the tests and a detailed overview of the implemented tests and corresponding test statistics can be found in the `vignette("robnptests")`.

Applications

Besides conventional two-sample problems, the tests can be applied in a moving time window for the online detection of structural breaks in outlier-contaminated time series. Abbas et al. (2016) describe how intensity changes in image sequences generated by a virus sensor are automatically detected by applying the tests to the individual pixel time series of the sequence. Moreover, the test statistics can be used as control statistics for robust, (approximately) distribution-free control charts for time series with a time-varying signal (Abbas & Fried, 2017, 2020). In Abbas et al. (2019), the tests are applied to detect unusual sequences in time series of crack widths in concrete beams by searching for sudden scale changes.

References

Abbas, S., & Fried, R. (2017). Control charts for the mean based on robust two-sample tests. *Journal of Statistical Computation and Simulation*, *87*(1), 138–155. https://doi.org/10.1080/00949655.2016.1194839

Abbas, S., & Fried, R. (2020). Robust control charts for the mean of a locally linear time series. *Journal of Statistical Computation and Simulation*, *90*(15), 2741–2765. https://doi.org/10.1080/00949655.2020.1788562

Abbas, S., Fried, R., & Gather, U. (2016). Detection of local intensity changes in grayscale images with robust methods for time-series analysis. In S. Michaelis, N. Piatkowski, & M. Stolpe (Eds.), *Solving large scale learning tasks: Challenges and algorithms: Essays dedicated to katharina morik on the occasion of her 60th birthday* (pp. 251–271). Springer. https://doi.org/10.1007/978-3-319-41706-6

Abbas, S., Fried, R., Heinrich, J., Horn, M., Jakubzik, M., Kohlenbach, J., Maurer, R., Michels, A., & Müller, C. H. (2019). Detection of anomalous sequences in crack data of a bridge monitoring. In N. Bauer, K. Ickstadt, K. Lübke, G. Szepannek, H. Trautmann, & M. Vichi (Eds.), *Applications in statistical computing: From music data analysis to industrial quality improvement* (pp. 251–269). Springer. https://doi.org/10.1007/978-3-030-25147-5_16

Fried, R. (2012). On the online estimation of piecewise constant volatilities. *Computational Statistics & Data Analysis*, *56*(11), 3080–3090. https://doi.org/10.1016/j.csda.2011.02.012

Fried, R., & Dehling, H. (2011). Robust nonparametric tests for the two sample location problem. *Statistical Methods & Applications*, *20*(4), 409–422. https://doi.org/10.1007/s10260-011-0164-1

Fried, R., & Gather, U. (2007). On rank tests for shift detection in time series. *Computational Statistics & Data Analysis*, *52*, 221–233. https://doi.org/10.1016/j.csda.2006.12.017

Helwig, N. E. (2021). *nptest: Nonparametric bootstrap and permutation tests*. https://CRAN.R-project.org/package=nptest

Hodges, J. L., & Lehmann, E. L. (1963). Estimates of location based on rank tests. *The Annals of Mathematical Statistics*, *34*(2), 598–611. https://doi.org/10.1214/aoms/1177704172

Huber, P. J. (1964). Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, *35*(1), 73–101. https://doi.org/10.1214/aoms/1177703732

Mair, P., & Wilcox, R. (2020). Robust statistical methods in R using the WRS2 package. *Behavior Research Methods*, *52*, 464–488. https://doi.org/10.3758/s13428-019-01246-w

Maronna, R. A., Martin, D. R., Yohai, V. J., & Salibián-Barrera, M. (2019). *Robust Statistics: Theory and Methods (with R)* (Second edition). Wiley. https://doi.org/10.1002/9781119214656

R Core Team. (2022). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.R-project.org/

Staudte, R. G., & Sheather, S. J. (1990). *Robust estimation and testing*. Wiley. https://doi.org/10.1002/9781118165485

Wilcox, R. R. (2003). *Applying contemporary statistical techniques*. Academic Press. https://doi.org/10.1016/B978-0-12-751541-0.X5021-4

Yuen, K. K., & Dixon, W. T. (1973). The approximate behaviour and performance of the two-sample trimmed t. *Biometrika*, *60*(2), 369–374. https://doi.org/10.2307/2334550