

Package ‘lookout’

January 19, 2026

Type Package

Title Leave One Out Kernel Density Estimates for Outlier Detection

Version 2.0.0

Maintainer Sevvandi Kandanaarachchi <sevvandik@gmail.com>

Description Outlier detection using leave-one-out kernel density estimates and extreme value theory. The bandwidth for kernel density estimates is computed using persistent homology, a technique in topological data analysis. Using peak-over-threshold method, a generalized Pareto distribution is fitted to the log of leave-one-out kde values to identify outliers.

License GPL-3

Encoding UTF-8

RoxygenNote 7.3.3

BugReports <https://github.com/sevvandi/lookout/issues>

Imports evd, ggplot2, RANN, robustbase, stats, TDAs, tidyverse

Suggests knitr, rmarkdown

URL <https://sevvandi.github.io/lookout/>,

<https://github.com/sevvandi/lookout>

NeedsCompilation no

Author Sevvandi Kandanaarachchi [aut, cre] (ORCID: <<https://orcid.org/0000-0002-0337-0395>>),
Rob Hyndman [aut] (ORCID: <<https://orcid.org/0000-0002-2140-5352>>),
Chris Fraley [ctb]

Repository CRAN

Date/Publication 2026-01-19 06:50:25 UTC

Contents

autplot.lookoutliers	2
autplot.persistingoutliers	3
find_tda_bw	4

lookout	5
lookout_ts	6
mvscale	7
persisting_outliers	9

Index	11
--------------	-----------

autoplot.lookoutliers *Plots outliers identified by lookout algorithm.*

Description

Scatterplot of two columns from the data set with outliers highlighted.

Usage

```
## S3 method for class 'lookoutliers'
autoplot(object, columns = 1:2, ...)
```

Arguments

object	The output of the function <code>lookout</code> .
columns	Which columns of the original data to plot (specified as either numbers or strings)
...	Other arguments currently ignored.

Value

A `ggplot` object.

Examples

```
X <- rbind(
  data.frame(
    x = rnorm(500),
    y = rnorm(500)
  ),
  data.frame(
    x = rnorm(5, mean = 10, sd = 0.2),
    y = rnorm(5, mean = 10, sd = 0.2)
  )
)
lo <- lookout(X)
autoplot(lo)
```

autoflot.persistingoutliers

Plots outlier persistence for a range of significance levels.

Description

This function plots outlier persistence for a range of significance levels using the algorithm lookout, an outlier detection method that uses leave-one-out kernel density estimates and generalized Pareto distributions to find outliers.

Usage

```
## S3 method for class 'persistingoutliers'  
autoflot(object, alpha = object$alpha, ...)
```

Arguments

object The output of the function `persisting_outliers`.
alpha The significance levels to plot.
... Other arguments currently ignored.

Value

A `ggplot` object.

Examples

```
X <- rbind(  
  data.frame(  
    x = rnorm(500),  
    y = rnorm(500)  
  ),  
  data.frame(  
    x = rnorm(5, mean = 10, sd = 0.2),  
    y = rnorm(5, mean = 10, sd = 0.2)  
  )  
)  
plot(X, pch = 19)  
outliers <- persisting_outliers(X, scale = FALSE)  
autoflot(outliers)
```

find_tda_bw	<i>Identifies bandwidth for outlier detection.</i>
-------------	--

Description

This function identifies the bandwidth that is used in the kernel density estimate computation. The function uses topological data analysis (TDA) to find the bandwidth.

Usage

```
find_tda_bw(X, fast = TRUE, gamma = 0.97, use_differences = FALSE)
```

Arguments

X	The numerical input data in a data.frame, matrix or tibble format.
fast	If TRUE (default), makes the computation faster by sub-setting the data for the bandwidth calculation.
gamma	Parameter for bandwidth calculation giving the quantile of the Rips death radii to use for the bandwidth. Default is 0.97. Ignored under the old version; where the lower limit of the maximum Rips death radii difference is used. Also ignored if bw is provided.
use_differences	If TRUE, the bandwidth is set to the lower point of the maximum Rips death radii differences. If FALSE, the gamma quantile of the Rips death radii is used. Default is FALSE.

Value

The bandwidth

Examples

```
X <- rbind(
  data.frame(
    x = rnorm(500),
    y = rnorm(500)
  ),
  data.frame(
    x = rnorm(5, mean = 10, sd = 0.2),
    y = rnorm(5, mean = 10, sd = 0.2)
  )
)
find_tda_bw(X, fast = TRUE)
```

lookout	<i>Identifies outliers using the algorithm lookout.</i>
---------	---

Description

This function identifies outliers using the algorithm lookout, an outlier detection method that uses leave-one-out kernel density estimates and generalized Pareto distributions to find outliers.

Usage

```
lookout(
  X,
  alpha = 0.01,
  beta = 0.9,
  gamma = 0.97,
  bw = NULL,
  gpd = NULL,
  scale = TRUE,
  fast = NROW(X) > 1000,
  old_version = FALSE
)
```

Arguments

X	The numerical input data in a data.frame, matrix or tibble format.
alpha	The level of significance. Default is 0.01 . So there is a 1/100 chance of any point being falsely classified as an outlier.
beta	The quantile threshold used in the GPD estimation. Default is 0.90 . To ensure there is enough data available, values greater than 0.90 are set to 0.90.
gamma	Parameter for bandwidth calculation giving the quantile of the Rips death radii to use for the bandwidth. Default is 0.97 . Ignored under the old version; where the lower limit of the maximum Rips death radii difference is used. Also ignored if bw is provided.
bw	Bandwidth parameter. If NULL (default), the bandwidth is found using Persistent Homology.
gpd	Generalized Pareto distribution parameters. If NULL (the default), these are estimated from the data.
scale	If TRUE, the data is standardized. Using the old version, unit scaling is applied so that each column is in the range $[0, 1]$. Under the new version, robust rotation and scaling is used so that the columns are approximately uncorrelated with unit variance. Default is TRUE.
fast	If TRUE (default), makes the computation faster by sub-setting the data for the bandwidth calculation.
old_version	Logical indicator of which version of the algorithm to use. Default is FALSE, meaning the newer version is used.

Value

A list with the following components:

outliers	The set of outliers.
outlier_probability	The GPD probability of the data.
outlier_scores	The outlier scores of the data.
bandwidth	The bandwidth selected using persistent homology.
kde	The kernel density estimate values.
lookde	The leave-one-out kde values.
gpd	The fitted GPD parameters.

References

Kandanaarachchi, S, and Hyndman, RJ (2022) Leave-one-out kernel density estimates for outlier detection, *J Computational & Graphical Statistics*, **31**(2), 586-599. <https://robjhyndman.com/publications/lookout/>.

Hyndman, RJ, Kandanaarachchi, S, and Turner, K (2026) When lookout meets crackle: Anomaly detection using kernel density estimation, in preparation. <https://robjhyndman.com/publications/lookout2.html>

Examples

```
X <- rbind(
  data.frame(
    x = rnorm(500),
    y = rnorm(500)
  ),
  data.frame(
    x = rnorm(5, mean = 10, sd = 0.2),
    y = rnorm(5, mean = 10, sd = 0.2)
  )
)
lo <- lookout(X)
lo
autoplot(lo)
```

lookout_ts	<i>Identifies outliers in univariate time series using the algorithm lookout.</i>
------------	---

Description

This is the time series implementation of lookout which identifies outliers in the double differenced time series.

Usage

```
lookout_ts(x, scale = FALSE, ...)
```

Arguments

x	The input univariate time series.
scale	If TRUE, the data is standardized. Using the old version, unit scaling is applied so that each column is in the range $[0, 1]$. Under the new version, robust rotation and scaling is used so that the columns are approximately uncorrelated with unit variance. Default is TRUE.
...	Other arguments are passed to lookout .

Value

A lookout object.

See Also

[lookout](#)

Examples

```
set.seed(1)
x <- arima.sim(list(order = c(1, 1, 0), ar = 0.8), n = 200)
x[50] <- x[50] + 10
plot(x)
lo <- lookout_ts(x)
lo
```

mvscale

Compute robust multivariate scaled data

Description

A multivariate version of [base:::scale\(\)](#), that takes account of the covariance matrix of the data, and uses robust estimates of center, scale and covariance by default. The centers are removed using medians, the scale function is the IQR, and the covariance matrix is estimated using a robust OGK estimate. The data are scaled using the Cholesky decomposition of the inverse covariance. Then the scaled data are returned.

Usage

```
mvscale(
  object,
  center = stats:::median,
  scale = robustbase:::s_Qn,
  cov = robustbase:::covOGK,
  warning = TRUE
)
```

Arguments

<code>object</code>	A vector, matrix, or data frame containing some numerical data.
<code>center</code>	A function to compute the center of each numerical variable. Set to <code>NULL</code> if no centering is required.
<code>scale</code>	A function to scale each numerical variable. When <code>cov = robustbase::covOGK()</code> , it is passed as the <code>sigmamu</code> argument.
<code>cov</code>	A function to compute the covariance matrix. Set to <code>NULL</code> if no rotation required.
<code>warning</code>	Should a warning be issued if non-numeric columns are ignored?

Details

Optionally, the centering and scaling can be done for each variable separately, so there is no rotation of the data, by setting `cov = NULL`. Also optionally, non-robust methods can be used by specifying `center = mean`, `scale = stats::sd()`, and `cov = stats::cov()`. Any non-numeric columns are retained with a warning.

Value

A vector, matrix or data frame of the same size and class as `object`, but with numerical variables replaced by scaled versions.

Author(s)

Rob J Hyndman

See Also

[base:::scale\(\)](#), [stats:::sd\(\)](#), [stats:::cov\(\)](#), [robustbase:::covOGK\(\)](#), [robustbase:::s_Qn\(\)](#)

Examples

```
# Univariate z-scores (no rotation)
z <- mvscale(faithful, center = mean, scale = sd, cov = NULL, warning = FALSE)
# Non-robust scaling with rotation
z <- mvscale(faithful, center = mean, cov = stats:::cov, warning = FALSE)
# Robust scaling and rotation
z <- mvscale(faithful, warning = FALSE)
```

`persisting_outliers` *Computes outlier persistence for a range of significance values.*

Description

This function computes outlier persistence for a range of significance values, using the algorithm lookout, an outlier detection method that uses leave-one-out kernel density estimates and generalized Pareto distributions to find outliers.

Usage

```
persisting_outliers(
  X,
  alpha = seq(0.01, 0.1, by = 0.01),
  st_qq = 0.9,
  scale = TRUE,
  num_steps = 20,
  old_version = FALSE
)
```

Arguments

<code>X</code>	The input data in a matrix, data.frame, or tibble format. All columns should be numeric.
<code>alpha</code>	Grid of significance levels.
<code>st_qq</code>	The starting quantile for death radii sequence. This will be used to compute the starting bandwidth value.
<code>scale</code>	If TRUE, the data is scaled. Default is TRUE. Which scaling method is used depends on the <code>old_version</code> parameter. See <code>lookout</code> for details.
<code>num_steps</code>	The length of the bandwidth sequence.
<code>old_version</code>	Logical indicator of which version of the algorithm to use.

Value

A list with the following components:

<code>out</code>	A 3D array of $N \times \text{num_steps} \times \text{num_alpha}$ where N denotes the number of observations, <code>num_steps</code> denote the length of the bandwidth sequence, and <code>num_alpha</code> denotes the number of significance levels. This is a binary array and the entries are set to 1 if that observation is an outlier for that particular bandwidth and significance level.
<code>bw</code>	The set of bandwidth values.
<code>gpdparas</code>	The GPD parameters used.
<code>lookoutbw</code>	The bandwidth chosen by the algorithm <code>lookout</code> using persistent homology.

Examples

```
X <- rbind(  
  data.frame(  
    x = rnorm(500),  
    y = rnorm(500)  
  ),  
  data.frame(  
    x = rnorm(5, mean = 10, sd = 0.2),  
    y = rnorm(5, mean = 10, sd = 0.2)  
  )  
)  
plot(X, pch = 19)  
outliers <- persisting_outliers(X, scale = FALSE)  
outliers  
autoplot(outliers)
```

Index

autoplot.lookoutliers, 2
autoplot.persistingoutliers, 3

base::scale(), 7, 8

find_tda_bw, 4

lookout, 5, 7, 9
lookout_ts, 6

mvscale, 7

persisting_outliers, 9

robustbase::covOGK(), 8
robustbase::s_Qn(), 8

stats::cov(), 8
stats::sd(), 8