

Package ‘outlierensembles’

July 22, 2025

Type Package

Title A Collection of Outlier Ensemble Algorithms

Version 0.1.3

Maintainer Sevvandi Kandanaarachchi <sevvandik@gmail.com>

Description Ensemble functions for outlier/anomaly detection. There is a new ensemble method proposed using Item Response Theory. Existing outlier ensemble methods from Schubert et al (2012) <[doi:10.1137/1.9781611972825.90](https://doi.org/10.1137/1.9781611972825.90)>, Chiang et al (2017) <[doi:10.1016/j.jal.2016.12.002](https://doi.org/10.1016/j.jal.2016.12.002)> and Aggarwal and Sathe (2015) <[doi:10.1145/2830544.2830549](https://doi.org/10.1145/2830544.2830549)> are also included.

License GPL (>= 3)

Encoding UTF-8

Depends R (>= 3.5)

Imports airt, EstCRM, psych, apcluster

RoxygenNote 7.3.2

Suggests dbscan, knitr, rmarkdown, ggplot2

VignetteBuilder knitr

URL <https://sevvandi.github.io/outlierensembles/>

LazyData true

NeedsCompilation no

Author Sevvandi Kandanaarachchi [aut, cre] (ORCID:
<<https://orcid.org/0000-0002-0337-0395>>)

Repository CRAN

Date/Publication 2025-03-27 04:10:01 UTC

Contents

average_ensemble	2
greedy_ensemble	3

icwa_ensemble	4
irt_ensemble	5
max_ensemble	6
threshold_ensemble	7
Y	8
Yfaithful	8
Index	9

average_ensemble	<i>Uses the mean as the ensemble score</i>
------------------	--

Description

This function uses the mean as the ensemble score.

Usage

```
average_ensemble(X)
```

Arguments

X The input data containing the outlier scores in a dataframe, matrix or tibble format. Rows contain observations and columns contain outlier detection methods.

Value

The ensemble scores.

Examples

```
set.seed(123)
if (requireNamespace("dbscan", quietly = TRUE)) {
X <- data.frame(x1 = rnorm(200), x2 = rnorm(200))
X[199, ] <- c(4, 4)
X[200, ] <- c(-3, 5)
# Using different parameters of lof for anomaly detection
y1 <- dbscan::lof(X, minPts = 10)
y2 <- dbscan::lof(X, minPts = 20)
knnobj <- dbscan::kNN(X, k = 20)
# Using different KNN distances as anomaly scores
y3 <- knnobj$dist[,10]
y4 <- knnobj$dist[,20]
# Dense points are less anomalous. Hence 1 - pointdensity is used.
y5 <- 1 - dbscan::pointdensity(X, eps = 0.8, type = "gaussian")
y6 <- 1 - dbscan::pointdensity(X, eps = 0.5, type = "gaussian")
Y <- cbind.data.frame(y1, y2, y3, y4, y5, y6)
ens <- average_ensemble(Y)
ens
}
```

greedy_ensemble	<i>Computes an ensemble score using the greedy algorithm proposed by Schubert et al (2012)</i>
-----------------	--

Description

This function computes an ensemble score using the greedy algorithm in the paper titled Evaluation of Outlier Rankings and Outlier Scores by Schubert et al (2012) <doi:10.1137/1.9781611972825.90>. The greedy ensemble is detailed in Section 4.3.

Usage

```
greedy_ensemble(X, kk = 5)
```

Arguments

X	The input data containing the outlier scores in a dataframe, matrix or tibble format. Rows contain observations and columns contain outlier detection methods.
kk	The number of estimated outliers.

Value

A list with the components:

scores	The ensemble scores.
methods	The methods that are chosen for the ensemble.
chosen	The chosen subset of original anomaly scores.

Examples

```
set.seed(123)
if (requireNamespace("dbscan", quietly = TRUE)){
  X <- data.frame(x1 = rnorm(200), x2 = rnorm(200))
  X[199, ] <- c(4, 4)
  X[200, ] <- c(-3, 5)
  # Using different parameters of lof for anomaly detection
  y1 <- dbscan::lof(X, minPts = 10)
  y2 <- dbscan::lof(X, minPts = 20)
  knnobj <- dbscan::kNN(X, k = 20)
  # Using different KNN distances as anomaly scores
  y3 <- knnobj$dist[,10]
  y4 <- knnobj$dist[,20]
  # Dense points are less anomalous. Hence 1 - pointdensity is used.
  y5 <- 1 - dbscan::pointdensity(X, eps = 0.8, type = "gaussian")
  y6 <- 1 - dbscan::pointdensity(X, eps = 0.5, type = "gaussian")
  Y <- cbind.data.frame(y1, y2, y3, y4, y5, y6)
  ens <- greedy_ensemble(Y, kk=5)
  ens$scores
```

```
}
```

icwa_ensemble	<i>Computes an ensemble score using inverse cluster weighted averaging method by Chiang et al (2017)</i>
---------------	--

Description

This function computes an ensemble score using inverse cluster weighted averaging in the paper titled A Study on Anomaly Detection Ensembles by Chiang et al (2017) <doi:10.1016/j.jal.2016.12.002>. The ensemble is detailed in Algorithm 2.

Usage

```
icwa_ensemble(X)
```

Arguments

X The input data containing the outlier scores in a dataframe, matrix or tibble format. Rows contain observations and columns contain outlier detection methods.

Value

The ensemble scores.

Examples

```
set.seed(123)
if (requireNamespace("dbscan", quietly = TRUE)) {
  X <- data.frame(x1 = rnorm(200), x2 = rnorm(200))
  X[199, ] <- c(4, 4)
  X[200, ] <- c(-3, 5)
  # Using different parameters of lof for anomaly detection
  y1 <- dbscan::lof(X, minPts = 10)
  y2 <- dbscan::lof(X, minPts = 20)
  knnobj <- dbscan::kNN(X, k = 20)
  # Using different KNN distances as anomaly scores
  y3 <- knnobj$dist[,10]
  y4 <- knnobj$dist[,20]
  # Dense points are less anomalous. Hence 1 - pointdensity is used.
  y5 <- 1 - dbscan::pointdensity(X, eps = 0.8, type = "gaussian")
  y6 <- 1 - dbscan::pointdensity(X, eps = 0.5, type = "gaussian")
  Y <- cbind.data.frame(y1, y2, y3, y4, y5, y6)
  ens <- icwa_ensemble(Y)
  ens
}
```

irt_ensemble

*Computes an ensemble score using Item Response Theory***Description**

This function computes an ensemble score using Item Response Theory (IRT). This was proposed as an ensemble method for anomaly/outlier detection in Kandanaarachchi (2021) <doi:10.13140/RG.2.2.18355.96801>.

Usage

```
irt_ensemble(X)
```

Arguments

X The input data containing the outlier scores in a dataframe, matrix or tibble format. Rows contain observations and columns contain outlier detection methods.

Details

For outlier detection, higher ensemble scores indicate higher levels of anomalousness. This ensemble uses IRT's latent trait to uncover the hidden ground truth, which is used as the ensemble score. It uses the R packages *airt* and *EstCRM* to fit the IRT models. It can also be used for other ensembling tasks.

Value

A list with the components:

scores The ensemble scores.
model The IRT model.

Examples

```
set.seed(123)
if (requireNamespace("dbscan", quietly = TRUE)) {
  X <- data.frame(x1 = rnorm(200), x2 = rnorm(200))
  X[199, ] <- c(4, 4)
  X[200, ] <- c(-3, 5)
  # Using different parameters of lof for anomaly detection
  y1 <- dbscan::lof(X, minPts = 10)
  y2 <- dbscan::lof(X, minPts = 20)
  knnobj <- dbscan::kNN(X, k = 20)
  # Using different KNN distances as anomaly scores
  y3 <- knnobj$dist[,10]
  y4 <- knnobj$dist[,20]
  # Dense points are less anomalous. Hence 1 - pointdensity is used.
  y5 <- 1 - dbscan::pointdensity(X, eps = 0.8, type = "gaussian")
  y6 <- 1 - dbscan::pointdensity(X, eps = 0.5, type = "gaussian")
  Y <- cbind.data.frame(y1, y2, y3, y4, y5, y6)
```

```

ens <- irt_ensemble(Y)
ens$scores
}

```

max_ensemble	<i>Computes an ensemble score using the maximum score of each observation</i>
--------------	---

Description

This function computes an ensemble score using the maximum score for each observation as detailed in Aggarwal and Sathe (2015) <doi:10.1145/2830544.2830549>.

Usage

```
max_ensemble(X)
```

Arguments

X	The input data containing the outlier scores in a dataframe, matrix or tibble format. Rows contain observations and columns contain outlier detection methods.
---	--

Value

The ensemble scores.

Examples

```

set.seed(123)
if (requireNamespace("dbscan", quietly = TRUE)) {
  X <- data.frame(x1 = rnorm(200), x2 = rnorm(200))
  X[199, ] <- c(4, 4)
  X[200, ] <- c(-3, 5)
  # Using different parameters of lof for anomaly detection
  y1 <- dbscan::lof(X, minPts = 10)
  y2 <- dbscan::lof(X, minPts = 20)
  knnobj <- dbscan::kNN(X, k = 20)
  # Using different KNN distances as anomaly scores
  y3 <- knnobj$dist[,10]
  y4 <- knnobj$dist[,20]
  # Dense points are less anomalous. Hence 1 - pointdensity is used.
  y5 <- 1 - dbscan::pointdensity(X, eps = 0.8, type = "gaussian")
  y6 <- 1 - dbscan::pointdensity(X, eps = 0.5, type = "gaussian")
  Y <- cbind.data.frame(y1, y2, y3, y4, y5, y6)
  ens <- max_ensemble(Y)
  ens
}

```

threshold_ensemble	<i>Computes an ensemble score by aggregating values above the mean</i>
--------------------	--

Description

This function computes an ensemble score by aggregating values above the mean as detailed in Aggarwal and Sathe (2015) <doi:10.1145/2830544.2830549>.

Usage

```
threshold_ensemble(X)
```

Arguments

X	The input data containing the outlier scores in a dataframe, matrix or tibble format. Rows contain observations and columns contain outlier detection methods.
---	--

Value

The ensemble scores.

Examples

```
set.seed(123)
if (requireNamespace("dbscan", quietly = TRUE)) {
  X <- data.frame(x1 = rnorm(200), x2 = rnorm(200))
  X[199, ] <- c(4, 4)
  X[200, ] <- c(-3, 5)
  # Using different parameters of lof for anomaly detection
  y1 <- dbscan::lof(X, minPts = 10)
  y2 <- dbscan::lof(X, minPts = 20)
  knnobj <- dbscan::kNN(X, k = 20)
  # Using different KNN distances as anomaly scores
  y3 <- knnobj$dist[,10]
  y4 <- knnobj$dist[,20]
  # Dense points are less anomalous. Hence 1 - pointdensity is used.
  y5 <- 1 - dbscan::pointdensity(X, eps = 0.8, type = "gaussian")
  y6 <- 1 - dbscan::pointdensity(X, eps = 0.5, type = "gaussian")
  Y <- cbind.data.frame(y1, y2, y3, y4, y5, y6)
  ens <- threshold_ensemble(Y)
  ens
}
```

Y	<i>A dataset containing anomaly scores</i>
---	--

Description

This dataset contains anomaly scores for the annulus data

Usage

Y

Format

A dataframe of 803 rows and 9 columns

Dimension 1 The different anomaly detection methods

Dimension 2 Each column has the anomaly scores from that method for the annulus dataset.

Yfaithful	<i>Old Faithful anomaly scores</i>
-----------	------------------------------------

Description

This dataset contains anomaly scores for the Old Faithful dataset

Usage

Yfaithful

Format

A dataframe of 272 rows and 9 columns

Dimension 1 The different anomaly detection methods

Dimension 2 Each column has the anomaly scores from that method for the old faithful dataset.

Index

* datasets

Y, [8](#)

Yfaithful, [8](#)

average_ensemble, [2](#)

greedy_ensemble, [3](#)

icwa_ensemble, [4](#)

irt_ensemble, [5](#)

max_ensemble, [6](#)

threshold_ensemble, [7](#)

Y, [8](#)

Yfaithful, [8](#)