

---

# **OutlierTree Documentation**

**David Cortes**

**May 30, 2022**



# CONTENTS

<b>1</b>	<b>Installation</b>	<b>3</b>
<b>2</b>	<b>Quick example notebook</b>	<b>5</b>
<b>3</b>	<b>Methods</b>	<b>7</b>
<b>4</b>	<b>OutlierTree</b>	<b>9</b>
<b>5</b>	<b>Indices and tables</b>	<b>15</b>
	<b>Bibliography</b>	<b>17</b>
	<b>Python Module Index</b>	<b>19</b>
	<b>Index</b>	<b>21</b>



This is the documentation page for the Python package *OutlierTree*, which does outlier/anomaly detection in tabular data with human-readable justifications. For more details, see the project's GitHub page:

<https://www.github.com/david-cortes/outliertree>

For the R version, see the CRAN webpage:

<https://cran.r-project.org/package=outliertree>



## INSTALLATION

Package is available on PyPI, can be installed with

```
pip install outlierstree
```



## QUICK EXAMPLE NOTEBOOK

- Explainable outlier detection in the Titanic dataset.



## METHODS

- *OutlierTree*
- *fit*
- *generate\_gritbot\_files*
- *predict*
- *print\_outliers*



## OUTLIERTREE

```
class outliertree.OutlierTree(max_depth=4, min_gain=0.01, z_norm=2.67, z_outlier=8.0,  
                             pct_outliers=0.01, min_size_numeric=25, min_size_categ=50,  
                             categ_split='binarize', categ_outliers='tail', numeric_split='raw',  
                             follow_all=False, gain_as_pct=True, nthreads=- 1)
```

Bases: object

Outlier Tree

Explainable outlier detection through decision-tree grouping. Tries to detect outliers by generating decision trees that attempt to “predict” the values of each column based on each other column, testing in each branch of every tried split (if it meets some minimum criteria) whether there are observations that seem too distant from the others in a 1-D distribution for the column that the split tries to “predict” (unlike other methods, this will not generate a score for each observation).

Splits are based on gain, while outlierness is based on confidence intervals. Similar in spirit to the GritBot software developed by RuleQuest research.

Supports columns of types numeric, categorical, and ordinal. Can handle missing values in any of them. Can also pass timestamps that will get converted to numeric but shown as timestamps in the output. Offers option to set columns to be used only for generating conditions without looking at outliers in them.

Infinite values will be taken into consideration when the column is used to split another column (that is,  $+\text{inf}$  will go into the branch that is greater than something,  $-\text{inf}$  into the other branch), but when a column is the target of the split, they will be taken as missing - that is, it will not report infinite values as outliers.

### Parameters

- **max\_depth** (*int*) – Maximum depth of the trees to grow. Can also pass zero, in which case it will only look for outliers with no conditions (i.e. takes each column as a 1-d distribution and looks for outliers in there independently of the values in other columns).
- **min\_gain** (*float*) – Minimum gain that a split has to produce in order to consider it (both in terms of looking for outliers in each branch, and in considering whether to continue branching from them). Note that default value for GritBot is  $1\text{e-}6$ , with ‘gain\_as\_pct’ = ‘False’, but it’s recommended to pass higher values (e.g.  $1\text{e-}1$ ) when using ‘gain\_as\_pct’ = ‘False’.
- **z\_norm** (*float*) – Maximum Z-value (from standard normal distribution) that can be considered as a normal observation. Note that simply having values above this will not automatically flag observations as outliers, nor does it assume that columns follow normal distributions. Also used for categorical and ordinal columns for building approximate confidence intervals of proportions.
- **z\_outlier** (*float*) – Minimum Z-value that can be considered as an outlier. There must be a large gap in the Z-value of the next observation in sorted order to consider it as outlier, given by  $(z\_outlier - z\_norm)$ . Decreasing this parameter is likely to result in more observations being flagged as outliers. Ignored for categorical and ordinal columns.

- **pct\_outliers** (*float*) – Approximate max percentage of outliers to expect in a given branch.
- **min\_size\_numeric** (*int*) – Minimum size that branches need to have when splitting a numeric column. In order to look for outliers in a given branch for a numeric column, it must have a minimum of twice this number of observations.
- **min\_size\_categ** (*int*) – Minimum size that branches need to have when splitting a categorical or ordinal column. In order to look for outliers in a given branch for a categorical, ordinal, or boolean column, it must have a minimum of twice this number of observations.
- **categ\_split** (*str*) – How to produce categorical-by-categorical splits. Options are:
  - "binarize"**:  
Will binarize the target variable according to whether it's equal to each present category within it (greater/less for ordinal), and split each binarized variable separately.
  - "bruteforce"**:  
Will evaluate each possible binary split of the categories (that is, it evaluates  $2^n$  potential splits every time). Note that trying this when there are many categories in a column will result in exponential computation time that might never finish.
  - "separate"**:  
Will create one branch per category of the splitting variable (this is how GritBot handles them).
- **categ\_outliers** (*str*) – How to look for outliers in categorical variables. Options are:
  - "tail"**:  
Will try to flag outliers if there is a large gap between proportions in sorted order, and this gap is unexpected given the prior probabilities. Such criteria tends to sometimes flag too many uninteresting outliers, but is able to detect more cases and recognize outliers when there is no single dominant category.
  - "majority"**:  
Will calculate an equivalent to z-value according to the number of observations that do not belong to the non-majority class, according to formula  $(n - n_{maj}) / (n * p_{prior}) < 1/z_{outlier}^2$ . Such criteria tends to miss many interesting outliers and will only be able to flag outliers in large sample sizes. This is the approach used by GritBot.
- **numeric\_split** (*str*) – How to determine the split point in numeric variables. Options are:
  - "mid"**:  
Will calculate the midpoint between the largest observation that goes to the '<=' branch and the smallest observation that goes to the '>' branch.
  - "raw"**:  
Will set the split point as the value of the largest observation that goes to the '<=' branch.

This doesn't affect how outliers are determined in the training data passed to 'fit', but it does affect the way in which they are presented and the way in which new outliers are detected when using 'predict'. "mid" is recommended for continuous-valued variables, while "raw" will provide more readable explanations for counts data at the expense of perhaps slightly worse generalizability to unseen data.
- **follow\_all** (*bool*) – Whether to continue branching from each split that meets the size and gain criteria. This will produce exponentially many more branches, and if depth is large, might take forever to finish. Will also produce a lot more spurious outliers. Not recommended.
- **gain\_as\_pct** (*bool*) – Whether the minimum gain above should be taken in absolute terms, or as a percentage of the standard deviation (for numerical columns) or shannon entropy (for categorical columns). Taking it in absolute terms will prefer making more splits on columns

that have a large variance, while taking it as a percentage might be more restrictive on them and might create deeper trees in some columns. For GritBot this parameter would always be 'False'. Recommended to pass higher values for 'min\_gain' when passing 'False' here. Not that when 'gain\_as\_pct' = 'False', the results will be sensitive to the scales of variables.

- **nthreads** (*int*) – Number of parallel threads to use. When fitting the model, it will only use up to one thread per column, while for prediction it will use up to one thread per row. The more threads that are used, the more memory will be required and allocated, so using more threads will not always lead to better speed. Negative numbers are interpreted as following joblib's formula (but not if the object attribute is overwritten after initialization). Can be changed after the object is already initialized.

### Variables

- **is\_fitted** (*bool*) – Indicates if the model as already been fit.
- **flaggable\_values** (*dict [ncols]*) – A dictionary indicating for each column which kind of values are possible to flag as outlier in at least one of the explored branches. For numerical and timestamp columns, will indicate the lower and upper bounds of the normal range (that is, it can only flag values *outside* of that interval), while for categorical, ordinal, and boolean, it will indicate the categories (which it can flag as outliers). If the lower bound is higher than the upper bound, it means that any value can potentially be flagged as outlier. If no values in a column can be flagged as outlier, entry for that column will be an empty dict.
- **cols\_num** (*array(ncols\_numeric, )*) – Names of the numeric columns in the data passed to '.fit'.
- **cols\_cat** (*array(ncols\_categ, )*) – Names of the categorical/string columns in the data passed to '.fit'.
- **cols\_bool** (*array(ncols\_bool, )*) – Names of the boolean columns in the data passed to '.fit'.
- **cols\_ord** (*array(ncols\_ordinal, )*) – Names of the ordinal columns in the data passed to '.fit'.
- **cols\_ts** (*array(ncols\_tiemstamp, )*) – Names of the timestamp columns in the data passed to '.fit'.

### References

**fit**(*df, cols\_ignore=None, outliers\_print=10, min\_decimals=2, return\_outliers=False*)

Fit Outlier Tree model to data.

---

**Note:** Row names will be taken as the index of the data frame. Column types will be taken as follows:

Numeric -> Everything that is a subtype of numpy's 'number' (e.g. integers and floats).

Categorical -> Python object (**even if the underlying types are numbers!!!**), boolean, and pandas Categorical.

Ordinal -> pandas Categorical with ordered attribute - the order will be the same as they have when passed.

Timestamp -> numpy's datetime64 dtype - they will only be taken with a precision of seconds due to numerical precision issues (code uses C 'long double'), and will be used as numerical but presented as timestamp in the outputs. If the required precision is less (e.g. just dates or years) it's recommended to pass them as numbers instead (e.g. 2008.1095 for some day in February 2008).

---

---

**Note:** You can look at the object’s attributes to ensure that the columns are being interpreted as the type (numeric, categorical, ordinal, boolean, timestamp) that they should have.

---

---

**Note:** Boolean or binary columns must be passed as categorical (will not accept them as numerical, nor as ordinal). If they have missing values, pandas will not be able to have them with dtype “bool” however (they need to be passed either as “object” or “Categorical” dtypes).

---

---

**Note:** Do NOT do one-hot or dummy encoding on categorical variables.

---

### Parameters

- **df** (*DataFrame(n\_rows, n\_cols)*) – Pandas’ DataFrame with regular (i.e. non-outlier) data that might contain some outliers.
- **cols\_ignore** (*boolean array(n\_cols, ) or string array(n\_ignore, )*) – Array containing columns which will not be split, but will be evaluated for usage in splitting other columns. Can pass either a boolean array with the same number of columns as ‘df’, or a list/array of column names (must match with those of ‘df’). Pass ‘None’ to use all columns.
- **outliers\_print** (*int or None*) – Maximum number of flagged outliers in the training data to print after fitting the model. Pass zero or None to avoid printing any. Outliers can be printed from resulting data frame afterwards through ‘.print\_outliers’.
- **min\_decimals** (*int*) – Minimum number of decimals to use when printing numeric values for the flagged outliers. The number of decimals will be dynamically increased according to the relative magnitudes of the values being reported. Ignored when passing `max_outliers=0` or `max_outliers=None`.
- **return\_outliers** (*bool*) – Whether to return a DataFrame with information about outliers flagged in the training data. If ‘True’, will return this information as a DataFrame. If ‘False’, will return this same object. See the documentation for ‘.predict’ for more information about the format.

### Returns

**result\_df or self** – Either a DataFrame with the information about potential outliers detected in the training data, or this same object if passing ‘return\_outliers’ = ‘False’. In the former, format is the same as when calling ‘.predict’. See the documentation for ‘.predict’ for more information about the output format.

### Return type

DataFrame(n\_rows, 6) or obj

**generate\_gritbot\_files**(*df, cols\_ignore=None, save\_folder='.', file\_name='data'*)

Generate data files for GritBot software

Generates CSV (.data) and naming (.names) files as required for use by the GritBot software (not included in this Python package). Note that this only generates the data files and doesn’t do anything else.

### Parameters

- **df** (*DataFrame(nrows, ncols)*) – Pandas DataFrame in the same format as required by ‘.fit’.

- **cols\_ignore** (*boolean array(n\_cols, ) or string array(n\_ignore, )*) – Array containing columns which will not be split, but will be evaluated for usage in splitting other columns. Pass ‘None’ to use all columns.
- **save\_folder** (*str*) – Path to folder where to save the generated files.
- **file\_name** (*str*) – Prefix for the file names (before the dot).

## References

**predict** (*df, outliers\_print=None, min\_decimals=2*)

Detect outliers on new data

---

**Note:** The group statistics for outliers are calculated only on rows that are not flagged as outliers - that is, they do not include the row being reported in their calculations, and exclude any outliers that were flagged at a previous parent branch.

---



---

**Note:** This will generate conditions having these criteria: “<=”, “>”, “in”, “=”, “!=” (not equal). When it says “in”, it means that the value is within the subset of categories provided. For ordinal columns, they will always be represented as “in”, but the subset will follow the order. The “.print” method will additionally simplify conditions to numeric “between”, and merge repeated splits that are on the same column.

---

## Parameters

- **df** (*DataFrame(n\_rows, n\_cols)*) – Pandas DataFrame in the same format and with the same columns as the one that was passed to ‘.fit’.
- **outliers\_print** (*int or None*) – Maximum number of outliers to print, if any are found. Pass zero or None to avoid printing any. Outliers can be printed from resulting data frame afterwards through ‘.print\_outliers’.
- **min\_decimals** (*int*) – Minimum number of decimals to use when printing numeric values for the flagged outliers. The number of decimals will be dynamically increased according to the relative magnitudes of the values being reported. Ignored when passing `max_outliers=0` or `max_outliers=None`.

## Returns

**result\_df** – DataFrame indicating for each row whether it is a suspected outlier or not. When they are suspected outliers, the resulting DataFrame will contain columns with

- information about the column in the input data whose value is suspicious,
- aggregate statistics for the values in this column among the normal observations,
- conditions that qualify the row to be put into that group,
- depth in the decision tree branch in which they can be considered an outlier,
- Whether the conditions that make it belong to the decision tree branch contain any NA split,
- The resulting outlier score or probability (based on Chebyshyov’s bound for numerical columns, and a simple upper confidence bound for categorical and ordinal columns - but note that it will always prefer to assign a row to an outlier branch that does not follow any NA path, or failing that, to the one with the smallest depth) - lower scores indicate lower probabilities and thus more outlierness.

### Return type

DataFrame(*n\_rows*, 6)

Information such as the conditions in the tree or the group statistics are returned as dictionaries, since they are not tabular format. The index of the output will be the same as that of the input. Outliers can be printed from resulting data frame afterwards through `‘.print_outliers’`.

**print\_outliers**(*df\_outliers*, *max\_outliers*=10, *min\_decimals*=2)

Print outliers in readable format

See the documentation for `‘predict’` for more details. This function will additionally perform some simplifications on the branch conditions, such as taking the smallest value when it is split two times by `“<=”`. Can also pass a smaller DataFrame with only selected outliers to be printed.

### Parameters

- **df\_outliers** (*DataFrame*(*n\_rows*, 6)) – DataFrame with outliers information as output by `‘fit’` or `‘predict’`.
- **max\_outliers** (*int*) – Maximum number of outliers to print.
- **min\_decimals** (*int*) – Minimum number of decimals to use when printing numeric values for the flagged outliers. The number of decimals will be dynamically increased according to the relative magnitudes of the values being reported. Ignored when passing `max_outliers=0`.

### Returns

**None** – No return value.

### Return type

None

## INDICES AND TABLES

- genindex
- modindex
- search



## BIBLIOGRAPHY

[gritbot] GritBot software : <https://www.rulequest.com/gritbot-info.html>



## PYTHON MODULE INDEX

### O

outliertree, 9



## INDEX

### F

`fit()` (*outliertree.OutlierTree* method), 11

### G

`generate_gritbot_files()` (*outliertree.OutlierTree* method), 12

### M

module  
    *outliertree*, 9

### O

*outliertree*  
    module, 9

*OutlierTree* (class in *outliertree*), 9

### P

`predict()` (*outliertree.OutlierTree* method), 13

`print_outliers()` (*outliertree.OutlierTree* method), 14