
Khiva Documentation

Release v0.2.0

Shapelets.io

Mar 11, 2019

Table of Contents

1	Getting Started	3
2	Khiva API	9
3	Bindings	49
4	Building programs using Khiva with CMake	51
5	AUTHORS	53
6	Cite Us	55

This is the documentation of **Khiva** library.

Khiva¹ is an open-source library of efficient algorithms to analyse time series in GPU and CPU. It can be used to extract insights from one or a group of time series. The large number of available methods allow us to understand the nature of each time series. Based on the results of this analysis, users can reduce dimensionality, find out recurrent motifs or discords, understand the seasonality or trend from a given time series, forecasting and detect anomalies.

Khiva provides a mean for time series analytics at scale. These analytics can be exploited in a wide range of use cases across several industries, like energy, finance, e-health, IoT, music industry, etc.

This is just the beginning, so stay tuned as more features are coming ...

[Gitter](#) is the place for discussions and questions about Khiva library. We use the [GitHub Issue Tracker](#) to manage bug reports and feature requests.

You can jump right into the package by looking into our [Getting Started](#).

¹ Time series analysis comprises methods for analyzing time series data in order to extract meaningful statistics and other characteristics of the data (Source [Wikipedia](#)).

1.1 Getting the source code

You can download the [latest stable released version](#), or you can get the latest source code version by cloning our git repository:

```
git clone https://github.com/shapelets/khiva
```

1.2 Dependencies

Khiva relies on a number of open source libraries and tools which are required to get it running.

Tools:

- A Build manager to control the compilation process [CMake](#).
- A library dependency manager [Conan](#).
- [Python 3](#).
- [Pip3](#).
- Documentation builders [Doxygen](#) and [sphinx](#).
- [Graphviz](#) and [Dot](#).
- A C++ compiler, it can be either [Clang](#), [GCC](#) or [Visual Studio C++ Compiler](#).

Note: All versions of Khiva Library require a **fully C++11-compliant** compiler.

Libraries:

- OpenCL library for you GPU card ([Intel](#), [AMD](#), or [Nvidia](#)).

- To run on accelerators like GPUs, [Arrayfire 3.6.2](#). Note that in order to use Arrayfire on Windows you need to [install](#) the Visual Studio 2015 (x64) runtime libraries.
- To test the functionality provided by Khiva, [Google Test](#).
- To benchmark Khiva, [Google Benchmark](#).
- [Boost](#).
- [Eigen](#).

1.3 Windows

1.3.1 Installation

Prerequisites

- Install [Python-64bits](#) and add the installation path to the environment variable path, 32-bits version won't work.
- Install [ArrayFire 3.6.2](#) and add the installation path to the environment variable path.
- Install [Vcpkg](#) and add the installation path to the environment variable path.
- Install [chocolatey](#) to manage windows dependencies and add the installation path to the environment variable path.

Once we have installed all Khiva dependencies, we are ready to install Khiva by using the installers (Option 1) or from source code (Option 2).

(Option 1) Build using a batch installer

In the tools directory you can find the script `install.bat`. You must indicate the path to your `vcpkg` installation directory.

- Usage: `install.bat <path_to_vcpkg>`
- Example: `install.bat c:\vcpkg`

(Option 2) Build from source code

If you prefer, you can build Khiva step by step. First, go to the source directory.

- Run `choco install cmake.install -NoNewWindow -Wait` Note: Add the installation path to the environment variable path and **before** than chocolatey environment variable path.
- Run `choco install doxygen.install -NoNewWindow -Wait`.
- Run `choco install graphviz -NoNewWindow -Wait`.
- Run `python -m pip install --upgrade pip`.
- Run `pip3 install sphinx breathe sphinx_rtd_theme`.
- Run `vcpkg install --triplet x64-windows gtest eigen3 benchmark boost`.
- Create a `build` folder in the root path of the project.
- Browse inside the `build` folder.

- Run `cmake .. -DCMAKE_TOOLCHAIN_FILE="<PATH_TO_VPKG>/scripts/buildsystems/vcpkg.cmake" -DKHIVA_USE_CONAN=OFF -G "Visual Studio 15 2017 Win64"` (Note: Replace <PATH_TO_VPKG> with your vcpkg installation path and do not forget to clean the build directory every time before running this command).
- Run `cmake --build . --config Release -- /m` to compile.

Install Khiva library

To install Khiva just execute the following command:

- Run `cmake -DBUILD_TYPE=Release -P cmake_install.cmake`.

1.3.2 Generating the Khiva installer

We use Cpack and NSIS to generate the installer.

Notes: Before generating the installer, the project must be built by following the steps in the previous Build from source code section. The generated package is stored in the build folder.

- Run `choco install nsis -NoNewWindow -Wait`.
- The installer can be generated running the command `cpack -G NSIS`.

Note: We use the `cpack` command from `cmake`, be aware `chocolatey` has another `cpack` command. If you cannot run the proper command, check out the path from `cmake` is placed before the path from `chocolatey` in the environment variable path.

1.3.3 Generating documentation

- Run `pip install sphinx` to install [Sphinx](#).
- Browse to the root path of the project.
- Run `sphinx-build.exe -b html doc/sphinx/source/ build/doc/html/`.

1.4 Linux

We use [Ubuntu 16.04 LTS](#) as our linux distribution example.

1.4.1 Prerequisites

- Install [Python-64bits](#) or run `apt-get install python3 python3-pip`, 32-bits version won't work.
- Download [ArrayFire 3.6.2](#).
- Create destination folder `sudo mkdir -p /opt/arrayfire`
- Install [ArrayFire](#) `sudo bash arrayfire/ArrayFire-v3.6.2_Linux_x86_64.sh --prefix=/opt/arrayfire --skip-license`

Once we have installed all Khiva dependencies, we are ready to install Khiva from source code or by using the installers.

1.4.2 Build from source code

First, go to the source directory.

```
conan remote add conan-mpusz https://api.bintray.com/conan/mpusz/conan-mpusz
mkdir build
cd build
conan install .. --build missing
cmake ..
make -j8
make install
```

It installs the library in `/usr/local/lib` and `/usr/local/include` folders.

In case ArrayFire is not installed in the default directory, it is required to add the `Arrayfire lib` folder to the `LD_LIBRARY_PATH` environment variable.

```
export LD_LIBRARY_PATH="/pathToArrayfire/arrayfire/lib:$LD_LIBRARY_PATH"
```

1.4.3 Install Khiva library from source code

- Run `make install`.

1.4.4 Generating the khiva installer

We use CPack to generate the installers from source code.

Notes: Before generating the installer the project should be built following the process explained in the previous `Build from source code` section. The generated package will be stored in the `build` folder.

For linux, either a deb or a rpm installer package can be generated. This could be done by running the command `cpack -G DEB` or `cpack -G RPM` respectively inside the build folder.

1.4.5 Generating documentation

We use `sphinx + doxygen` to generate our documentation. You need to install the following packages:

- Sphinx: `brew install sphinx`.
- Doxygen: `brew install doxygen`.
- Read the Docs Theme: `pip install sphinx_rtd_theme`.
- Breathe: `pip install breathe`.

To generate the khiva documentation run the following command.

- Run `make documentation`.

1.5 Mac OS

1.5.1 Prerequisites

- Install [Python-64bits](#) or just run `brew install python3`, 32-bits version won't work.

- Install [ArrayFire 3.6.2](#) and then execute the following lines to move the ArrayFire files from the default installation directory to the system path for libraries:

```
sudo mv /opt/arrayfire/include/* /usr/local/include
sudo mv /opt/arrayfire/lib/* /usr/local/lib
sudo mv /opt/arrayfire/share/* /usr/local/share
sudo rm -rf /opt/arrayfire
```

Once we have installed all Khiva dependencies, we are ready to build and install Khiva, either by using the installers or from source code.

1.5.2 Build from source code

First, go to the directory where the source code is stored:

```
conan remote add conan-mpusz https://api.bintray.com/conan/mpusz/conan-mpusz
mkdir build
cd build
conan install .. --build missing
cmake ..
make -j8
make install
```

It installs the library in `/usr/local/lib` and `/usr/local/include` folders.

1.5.3 Install Khiva library from source code

- Run `make install`.

1.5.4 Generating the khiva installer

For Mac OS, the installer can be generated by running the command `cpack -G productbuild` inside the build folder. Note that, before generating the installer you have to follow the previous Build from source code section.

1.5.5 Generating documentation

We use `sphinx + doxygen` to generate our documentation. You will need to install the following packages:

- Sphinx: `brew install sphinx`.
- Doxygen: `brew install doxygen`.
- Read the Docs Theme: `pip install sphinx_rtd_theme`.
- Breathe: `pip install breathe`.

To generate the khiva documentation run the following command.

- `make documentation`.

This is the list of namespaces that comprise the Khiva library.

2.1 Namespace Array

namespace array

Functions

af::array **createArray** (void **data*, unsigned *ndims*, dim_t **dims*, **const** int *type*)
Creates an af::array.

Return af::array Containing the data.

Parameters

- *data*: Data used to create the af::array.
- *ndims*: Number of dimensions of data.
- *dims*: Cardinality of dimensions of data.
- *type*: Data type.

void **getData** (af::array *array*, void **data*)
Retrieves the data from the device to the host.

Parameters

- *array*: The [Array](#) that contains the data to be retrieved.
- *data*: Pointer to a preallocated block of memory in the host.

af::dim4 **getDims** (af::array *array*)
Returns the dimensions from a given array.

Return af::dim4 The dimensions.

Parameters

- array: [Array](#) from which to get the dimensions.

void **print** (af::array *array*)
Prints the content of an array.

Parameters

- array: The array to be printed.

void **deleteArray** (af_array *array*)
Decreases the references count for the given array.

Parameters

- array: The [Array](#) to be deleted.

int **getType** (af::array *array*)
Gets the type of the array.

Return int Value of the Dtype enumeration.

Parameters

- array: The array to obtain the type from.

template <typename T>
std::vector<int> **getRowsWithMaximals** (khiva::array::Array<T> *a*)
Gets the indices of all rows containing a maximal.

Return std::vector<int> with the indices of the rows with maximals.

Parameters

- a: The input array.

template <typename T>
std::vector<int> **getIndexMaxColumns** (std::vector<T> *r*)
Gets the indices of the columns with maximals.

Return std::vector<int> with the indices of the columns with maximals.

Parameters

- r: The input row.

template <class T>
class Array
#include </home/docs/checkouts/readthedocs.org/user_builds/khiva/checkouts/v0.2.0/include/khiva/array.h>
[Array](#) class, This class provides functionality manage Arrays on the host side.

Public Functions

Array()

Default Constructor of *Array* class.

Array(af::array in)

Constructor of *Array* class which receives and af::array.

Parameters

- in: The input af::array.

~Array()

Default destructor of *Array* class.

void **setNumX**(int val)

Sets the cardinality of the first dimension.

Parameters

- val: The value to be set.

void **setNumY**(int val)

Sets the cardinality of the second dimension.

Parameters

- val: The value to be set.

void **setNumW**(int val)

Sets the cardinality of the third dimension.

Parameters

- val: The value to be set.

void **setNumZ**(int val)

Sets the cardinality of the fourth dimension.

Parameters

- val: The value to be set.

void **setData**(T *pd)

Sets the data to be stored in the *Array*.

Parameters

- pd: The data to be stored.

int **getNumX**()

Gets the cardinality of the first dimension.

Return int the Cardinality of the first dimension.

int **getNumY**()

Gets the cardinality of the second dimension.

Return int the Cardinality of the second dimension.

int **getNumW**()

Gets the cardinality of the third dimension.

Return int the Cardinality of the third dimension.

int **getNumZ**()

Gets the cardinality of the fourth dimension.

Return int the Cardinality of the fourth dimension.

int **getNumElements** ()

Gets the number of elements in data.

Return int the Cardinality of the number of elements.

std::vector<T> **getRow** (int *idx*)

Gets the row number given by *idx*.

Return std::vector Containing the selected row.

Parameters

- *idx*: The row number to be extracted.

std::vector<T> **getColumn** (int *idx*)

Gets the column number given by *idx*.

Return std::vector Containing the selected column.

Parameters

- *idx*: The column number to be extracted.

T **getElement** (int *row*, int *column*)

Gets the element given by row and column.

Return T The element to be extracted.

Parameters

- *row*: The row number.
- *column*: The column number.

T* **getData** ()

Gets a pointer to the data stored in the array.

Return T Pointer to data.

bool **isEmpty** ()

Checks whether The *Array* is empty or not.

Return True if the *Array* is empty, false otherwise.

void **print** ()

Prints the content of the array.

2.2 Namespace Clustering

namespace clustering

Functions

void **kMeans** (af::array *tss*, int *k*, af::array &*centroids*, af::array &*labels*, float *tolerance* = 0.0000000001, int *maxIterations* = 100)
Calculates the k-means algorithm.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

- `k`: The number of means to be computed.
- `centroids`: The resulting means or centroids.
- `labels`: The resulting labels of each time series which is the closest centroid.
- `tolerance`: The error tolerance to stop the computation of the centroids.
- `maxIterations`: The maximum number of iterations allowed.

void **kShape** (af::array *tss*, int *k*, af::array &*centroids*, af::array &*labels*, float *tolerance* = 0.0000000001, int *maxIterations* = 100)
Calculates the k-shape algorithm.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `k`: The number of means to be computed.
- `centroids`: The resulting means or centroids.
- `labels`: The resulting labels of each time series which is the closest centroid.
- `tolerance`: The error tolerance to stop the computation of the centroids.
- `maxIterations`: The maximum number of iterations allowed.

2.3 Namespace Dimensionality

namespace dimensionality

Typedefs

typedef std::pair<float, float> **Point**

typedef std::pair<int, int> **Segment**

Functions

std::vector<[Point](#)> **PAA** (std::vector<[Point](#)> *points*, int *bins*)

Piecewise Aggregate Approximation (PAA) approximates a time series X of length n into vector $\bar{X} = (\bar{x}_1, \dots, \bar{x}_M)$ of any arbitrary length $M \leq n$ where each of \bar{x}_i is calculated as follows:

$$\bar{x}_i = \frac{M}{n} \sum_{j=n/M(i-1)+1}^{(n/M)i} x_j.$$

Which simply means that in order to reduce the dimensionality from n to M , we first divide the original time series into M equally sized frames and secondly compute the mean values for each frame. The sequence assembled from the mean values is the PAA approximation (i.e., transform) of the original time series.

Return result A vector of Points with the reduced dimensionality.

Parameters

- `points`: Set of points.
- `bins`: Sets the total number of divisions.

`af::array PAA (af::array a, int bins)`

Piecewise Aggregate Approximation (PAA) approximates a time series X of length n into vector $\bar{X} = (\bar{x}_1, \dots, \bar{x}_M)$ of any arbitrary length $M \leq n$ where each of \bar{x}_i is calculated as follows:

$$\bar{x}_i = \frac{M}{n} \sum_{j=n/M(i-1)+1}^{(n/M)i} x_j.$$

Which simply means that in order to reduce the dimensionality from n to M , we first divide the original time series into M equally sized frames and secondly compute the mean values for each frame. The sequence assembled from the mean values is the PAA approximation (i.e., transform) of the original time series.

Return `af::array` An array of points with the reduced dimensionality.

Parameters

- `a`: Set of points.
- `bins`: Sets the total number of divisions.

`af::array PIP (af::array ts, int numberIPs)`

Calculates the number of Perceptually Important Points (PIP) in the time series.

[1] Fu TC, Chung FL, Luk R, and Ng CM. Representing financial time series based on data point importance. *Engineering Applications of Artificial Intelligence*, 21(2):277-300, 2008.

Return `af::array` Array with the most Perceptually Important numPoints.

Parameters

- `ts`: Expects an input array whose dimension zero is the length of the time series.
- `numberIPs`: The number of points to be returned.

`std::vector<Point> PLABottomUp (std::vector<Point> ts, float maxError)`

Applies the Piecewise Linear Approximation (PLA BottomUP) to the time series.

[1] Zhu Y, Wu D, Li Sh (2007). A Piecewise Linear Representation Method of Time Series Based on Feature Points. *Knowledge-Based Intelligent Information and Engineering Systems* 4693:1066-1072.

Return `std::vector` Vector with the reduced number of points.

Parameters

- `ts`: Expects an input vector containing the set of points to be reduced.
- `maxError`: The maximum approximation error allowed.

`af::array PLABottomUp (af::array ts, float maxError)`

Applies the Piecewise Linear Approximation (PLA BottomUP) to the time series.

[1] Zhu Y, Wu D, Li Sh (2007). A Piecewise Linear Representation Method of Time Series Based on Feature Points. *Knowledge-Based Intelligent Information and Engineering Systems* 4693:1066-1072.

Return `af::array` with the reduced number of points.

Parameters

- `ts`: Expects an `af::array` containing the set of points to be reduced. The first component of the points in the first column and the second component of the points in the second column.
- `maxError`: The maximum approximation error allowed.

`std::vector<Point> PLASlidingWindow (std::vector<Point> ts, float maxError)`

Applies the Piecewise Linear Approximation (PLA Sliding Window) to the time series.

[1] Zhu Y, Wu D, Li Sh (2007). A Piecewise Linear Representation Method of Time Series Based on Feature Points. Knowledge-Based Intelligent Information and Engineering Systems 4693:1066-1072.

Return `std::vector` Vector with the reduced number of points.

Parameters

- `ts`: Expects an input vector containing the set of points to be reduced.
- `maxError`: The maximum approximation error allowed.

`af::array PLASlidingWindow (af::array ts, float maxError)`

Applies the Piecewise Linear Approximation (PLA Sliding Window) to the time series.

[1] Zhu Y, Wu D, Li Sh (2007). A Piecewise Linear Representation Method of Time Series Based on Feature Points. Knowledge-Based Intelligent Information and Engineering Systems 4693:1066-1072.

Return `af::array` with the reduced number of points.

Parameters

- `ts`: Expects an `af::array` containing the set of points to be reduced. The first component of the points in the first column and the second component of the points in the second column.
- `maxError`: The maximum approximation error allowed.

`std::vector<Point> ramerDouglasPeucker (std::vector<Point> pointList, double epsilon)`

The Ramer–Douglas–Peucker algorithm (RDP) is an algorithm for reducing the number of points in a curve that is approximated by a series of points. It reduces a set of points depending on the perpendicular distance of the points and epsilon, the greater epsilon, more points are deleted.

[1] Urs Ramer, “An iterative procedure for the polygonal approximation of plane curves”, Computer Graphics and Image Processing, 1(3), 244–256 (1972) doi:10.1016/S0146-664X(72)80017-0.

[2] David Douglas & Thomas Peucker, “Algorithms for the reduction of the number of points required to represent a

digitized line or its caricature”, The Canadian Cartographer 10(2), 112–122 (1973) doi:10.3138/FM57-6770-U75U-7727

Return `std::vector<khiva::dimensionality::Point>` with the selected points.

Parameters

- `pointList`: Set of input points.
- `epsilon`: It acts as the threshold value to decide which points should be considered meaningful or not.

af::array **ramerDouglasPeucker** (af::array *pointList*, double *epsilon*)

The Ramer–Douglas–Peucker algorithm (RDP) is an algorithm for reducing the number of points in a curve that is approximated by a series of points. It reduces a set of points depending on the perpendicular distance of the points and epsilon, the greater epsilon, more points are deleted.

[1] Urs Ramer, “An iterative procedure for the polygonal approximation of plane curves”, Computer Graphics and Image Processing, 1(3), 244–256 (1972) doi:10.1016/S0146-664X(72)80017-0.

[2] David Douglas & Thomas Peucker, “Algorithms for the reduction of the number of points required to represent a

digitized line or its caricature”, The Canadian Cartographer 10(2), 112–122 (1973) doi:10.3138/FM57-6770-U75U-7727

Return af::array with the selected points.

Parameters

- *pointList*: Set of input points.
- *epsilon*: It acts as the threshold value to decide which points should be considered meaningful or not.

af::array **SAX** (af::array *a*, int *alphabetSize*)

Symbolic Aggregate approXimation (SAX). It transforms a numeric time series into a time series of symbols with the same size. The algorithm was proposed by Lin et al.) and extends the PAA-based approach inheriting the original algorithm simplicity and low computational complexity while providing satisfactory sensitivity and selectivity in range query processing. Moreover, the use of a symbolic representation opened a door to the existing wealth of data-structures and string-manipulation algorithms in computer science such as hashing, regular expression, pattern matching, suffix trees, and grammatical inference.

[1] Lin, J., Keogh, E., Lonardi, S. & Chiu, B. (2003) A Symbolic Representation of Time Series, with Implications for Streaming Algorithms. In proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery. San Diego, CA. June 13.

Return result An array of symbols.

Parameters

- *a*: Array with the input time series.
- *alphabetSize*: Number of element within the alphabet.

std::vector<*Point*> **visvalingam** (std::vector<*Point*> *pointList*, int *numPoints*)

Reduces a set of points by applying the Visvalingam method (minimum triangle area) until the number of points is reduced to numPoints.

[1] M. Visvalingam and J. D. Whyatt, Line generalisation by repeated elimination of points, The Cartographic Journal, 1993.

Return std::vector<khiva::dimensionality::Point> where the number of points has been reduced to numPoints.

Parameters

- *pointList*: Expects an input vector of points.
- *numPoints*: Sets the number of points returned after the execution of the method.

af::array **visvalingam** (af::array *pointList*, int *numPoints*)

Reduces a set of points by applying the Visvalingam method (minimum triangle area) until the number of points is reduced to numPoints.

[1] M. Visvalingam and J. D. Whyatt, Line generalisation by repeated elimination of points, The Cartographic Journal, 1993.

Return af::array where the number of points has been reduced to numPoints.

Parameters

- *pointList*: Expects an input vector of points.
- *numPoints*: Sets the number of points returned after the execution of the method.

2.4 Namespace Distances

namespace distances

Functions

double **dtw** (std::vector<double> *a*, std::vector<double> *b*)

Calculates the Dynamic Time Warping Distance.

Return array The resulting distance between a and b.

Parameters

- *a*: The input time series of reference.
- *b*: The input query.

af::array **dtw** (af::array *tss*)

Calculates the Dynamic Time Warping Distance.

Return af::array An upper triangular matrix where each position corresponds to the distance between two time series. Diagonal elements will be zero. For example: Position row 0 column 1 records the distance between time series 0 and time series 1.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **euclidean** (af::array *tss*)

Calculates euclidean distances between time series.

Return af::array An upper triangular matrix where each position corresponds to the distance between two time series. Diagonal elements will be zero. For example: Position row 0 column 1 records the distance between time series 0 and time series 1.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **hamming** (af::array *tss*)

Calculates hamming distances between time series.

Return af::array An upper triangular matrix where each position corresponds to the distance between two time series. Diagonal elements will be zero. For example: Position row 0 column 1 records the distance between time series 0 and time series 1.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **manhattan** (af::array *tss*)

Calculates manhattan distances between time series.

Return af::array An upper triangular matrix where each position corresponds to the distance between two time series. Diagonal elements will be zero. For example: Position row 0 column 1 records the distance between time series 0 and time series 1.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **sbd** (af::array *tss*)

Calculates the Shape-Based distance (SBD). It computes the normalized cross-correlation and it returns 1.0 minus the value that maximizes the correlation value between each pair of time series.

Return array An upper triangular matrix where each position corresponds to the distance between two time series. Diagonal elements will be zero. For example: Position row 0 column 1 records the distance between time series 0 and time series 1.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **squaredEuclidean** (af::array *tss*)

Calculates non squared version of the euclidean distance.

Return array An upper triangular matrix where each position corresponds to the distance between two time series. Diagonal elements will be zero. For example: Position row 0 column 1 records the distance between time series 0 and time series 1.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

2.5 Namespace Features

namespace features

Functions

af::array **absEnergy** (af::array *base*)

Calculates the absolute energy of the time series which is the sum over the squared values.

$$E = \sum_{i=1, \dots, n} x_i^2$$

.

Return af::array An array with the same dimensions as *tss*, whose values (time series in dimension 0) contains the sum of the squares values in the time series.

Parameters

- *base*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **absoluteSumOfChanges** (af::array *tss*)

Calculates the sum over the absolute value of consecutive changes in the time series.

$$\sum_{i=1, \dots, n-1} |x_{i+1} - x_i|$$

.

Return af::array An array with the same dimensions as *tss*, whose values (time series in dimension 0) contains absolute value of consecutive changes in the time series.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **aggregatedAutocorrelation** (af::array *tss*, af::array (**aggregationFunction*)) **const**

af::array&, **const** bool, **const** dim_t

Calculates the value of an aggregation function *f_agg* (e.g. var or mean) of the autocorrelation (Compare to <http://en.wikipedia.org/wiki/Autocorrelation#Estimation>), taken over different all possible lags (1 to length of *x*).

$$\frac{1}{n-1} \sum_{l=1, \dots, n} \frac{1}{(n-l)\sigma^2} \sum_{t=1}^{n-l} (X_t - \mu)(X_{t+l} - \mu),$$

where *n* is the length of the time series X_i , σ^2 its variance and μ its mean.

Return af::array An array with the same dimensions as *tss*, whose values (time series in dimension 0) contains the aggregated correlation for each time series.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *aggregationFunction*: The function to summarise all autocorrelation with different lags.

af::array **aggregatedAutocorrelation** (af::array *tss*, af::array (**aggregationFunction*)) **const**
af::array&, **const** int

Calculates the value of an aggregation function *f_agg* (e.g. var or mean) of the autocorrelation (Compare to <http://en.wikipedia.org/wiki/Autocorrelation#Estimation>), taken over different all possible lags (1 to length of *x*).

$$\frac{1}{n-1} \sum_{l=1, \dots, n} \frac{1}{(n-l)\sigma^2} \sum_{t=1}^{n-l} (X_t - \mu)(X_{t+l} - \mu),$$

where *n* is the length of the time series *X_i*, σ^2 its variance and μ its mean.

Return af::array An array with the same dimensions as *tss*, whose values (time series in dimension 0) contains the aggregated correlation for each time series.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *aggregationFunction*: The function to summarise all autocorrelation with different lags.

af::array **aggregatedAutocorrelation** (af::array *tss*, af::array (**aggregationFunction*)) **const**
af::array&, **const** dim_t

Calculates the value of an aggregation function *f_agg* (e.g. var or mean) of the autocorrelation (Compare to <http://en.wikipedia.org/wiki/Autocorrelation#Estimation>), taken over different all possible lags (1 to length of *x*).

$$\frac{1}{n-1} \sum_{l=1, \dots, n} \frac{1}{(n-l)\sigma^2} \sum_{t=1}^{n-l} (X_t - \mu)(X_{t+l} - \mu),$$

where *n* is the length of the time series *X_i*, σ^2 its variance and μ its mean.

Return af::array An array with the same dimensions as *tss*, whose values (time series in dimension 0) contains the aggregated correlation for each time series.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *aggregationFunction*: The function to summarise all autocorrelation with different lags.

void **aggregatedLinearTrend** (af::array *t*, long *chunkSize*, af::array (**aggregationFunction*)) **const** af::array&, **const** int
, af::array &*slope*, af::array &*intercept*, af::array &*rvalue*, af::array &*pvalue*, af::array &*stderr*
restCalculates a linear least-squares regression for values of the time series that were aggregated over chunks versus the sequence from 0 up to the number of chunks minus one.

Parameters

- *t*: The time series to calculate the features of.
- *chunkSize*: The *chunkSize* used to aggregate the data.
- *aggregationFunction*: Function to be used in the aggregation.
- *slope*: Slope of the regression line.
- *intercept*: Intercept of the regression line.

- `rvalue`: Correlation coefficient.
- `pvalue`: Two-sided p-value for a hypothesis test whose null hypothesis is that the slope is zero, using Wald Test with t-distribution of the test statistic.
- `stderrest`: Standard error of the estimated gradient.

void **aggregatedLinearTrend** (af::array *t*, long *chunkSize*, af::array (**aggregationFunction*)) **const** af::array&, **const** dim_t, af::array &*slope*, af::array &*intercept*, af::array &*rvalue*, af::array &*pvalue*, af::array &*stderrest*) Calculates a linear least-squares regression for values of the time series that were aggregated over chunks versus the sequence from 0 up to the number of chunks minus one.

Parameters

- `t`: The time series to calculate the features of.
- `chunkSize`: The chunkSize used to aggregate the data.
- `aggregationFunction`: Function to be used in the aggregation.
- `slope`: Slope of the regression line.
- `intercept`: Intercept of the regression line.
- `rvalue`: Correlation coefficient.
- `pvalue`: Two-sided p-value for a hypothesis test whose null hypothesis is that the slope is zero, using Wald Test with t-distribution of the test statistic.
- `stderrest`: Standard error of the estimated gradient.

af::array **approximateEntropy** (af::array *tss*, int *m*, float *r*)

Calculates a vectorized Approximate entropy algorithm (https://en.wikipedia.org/wiki/Approximate_entropy). For short time series, this method is highly dependent on the parameters, but should be stable for $N > 2000$, see:

[1] Yentes et al., The Appropriate Use of Approximate Entropy and Sample Entropy with Short Data Sets, (2012). Other shortcomings and alternatives discussed in: Richman & Moorman, Physiological time-series analysis using approximate entropy and sample entropy, (2000).

Return af::array An array with the same dimensions as *tss*, whose values (time series in dimension 0) contains the vectorized Approximate entropy for all the input time series in *tss*.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `m`: Length of compared run of data.
- `r`: Filtering level, must be positive.

af::array **autoCorrelation** (af::array *tss*, long *maxLag*, bool *unbiased* = false)

Calculates the autocorrelation of the specified lag for the given time series, according to the formula [1].

$$\frac{1}{(n-l)\sigma^2} \sum_{t=1}^{n-l} (X_t - \mu)(X_{t+l} - \mu),$$

where n is the length of the time series X_i , σ^2 its variance and μ its mean, l denotes the lag.

[1] <https://en.wikipedia.org/wiki/Autocorrelation#Estimation>

Return af::array The autocorrelation value for the given time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `maxLag`: The maximum lag to compute.
- `unbiased`: Determines whether it divides by $(n - \text{lag})$ (if true), or n (if false).

af::array **autoCovariance** (af::array *xss*, bool *unbiased* = false)

Calculates the auto-covariance the given time series.

Return af::array The auto-covariance value for the given time series.

Parameters

- `xss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `unbiased`: Determines whether it divides by $n - \text{lag}$ (if true) or n (if false).

af::array **binnedEntropy** (af::array *tss*, int *max_bins*)

Calculates the binned entropy for the given time series and number of bins. It calculates the value of:

$$\sum_{k=0}^{\min(\text{max_bins}, \text{len}(x))} p_k \log(p_k) \cdot \mathbf{1}_{(p_k > 0)},$$

where p_k is the percentage of samples in bin k .

Return af::array The binned entropy value for the given time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `max_bins`: The number of bins.

af::array **c3** (af::array *tss*, long *lag*)

This function calculates the value of:

$$\frac{1}{n - 2\text{lag}} \sum_{i=0}^{n-2\text{lag}} x_{i+2\cdot\text{lag}}^2 \cdot x_{i+\text{lag}} \cdot x_i,$$

which is:

$$\mathbb{E}[L^2(X)^2 \cdot L(X) \cdot X],$$

where \mathbb{E} is the mean and L is the lag operator. It was proposed in [1] as a measure of non linearity in the time series.

[1] Schreiber, T. and Schmitz, A., Discrimination power of measures for nonlinearity in a time series, PHYSICAL REVIEW E, VOLUME 55, NUMBER 5, (1997).

Return af::array The non-linearity value for the given time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `lag`: The lag.

`af::array` **cidCe** (`af::array tss`, `bool zNormalize = false`)

This function calculator is an estimate for a time series complexity [1]. It calculates the value of:

$$\sqrt{\sum_{i=0}^{n-2lag} (x_i - x_{i+1})^2}.$$

[1] Batista, Gustavo EAPA, et al (2014). CID: an efficient complexity-invariant distance for time series. Data Mining and Knowledge Difscovery 28.3 (2014): 634-669.

Return `af::array` The complexity value for the given time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `zNormalize`: Controls whether the time series should be z-normalized or not.

`af::array` **countAboveMean** (`af::array tss`)

Calculates the number of values in the time series that are higher than the mean.

Return `af::array` The number of values in the time series that are higher than the mean.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **countBelowMean** (`af::array tss`)

Calculates the number of values in the time series that are lower than the mean.

Return `af::array` The number of values in the time series that are lower than the mean.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **crossCovariance** (`af::array xss`, `af::array yss`, `bool unbiased = true`)

Calculates the cross-covariance of the given time series.

Return `af::array` The cross-covariance value for the given time series.

Parameters

- `xss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `yss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `unbiased`: Determines whether it divides by `n - lag` (if true) or `n` (if false).

af::array **crossCorrelation** (af::array *xss*, af::array *yss*, bool *unbiased* = true)

Calculates the cross-correlation of the given time series.

Return af::array The cross-correlation value for the given time series.

Parameters

- *xss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *yss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *unbiased*: Determines whether it divides by $n - \text{lag}$ (if true) or n (if false).

af::array **cwtCoefficients** (af::array *tss*, af::array *widths*, int *coeff*, int *w*)

Calculates a Continuous wavelet transform for the Ricker wavelet, also known as the “Mexican hat wavelet” which is defined by:

$$\frac{2}{\sqrt{3a}\pi^{\frac{1}{4}}} \left(1 - \frac{x^2}{a^2}\right) \exp\left(-\frac{x^2}{2a^2}\right),$$

where a is the width parameter of the wavelet function. This feature calculator takes three different parameter: *widths*, *coeff* and *w*. The feature calculator takes all the different widths arrays and then calculates the cwt one time for each different width array. Then the values for the different coefficient for *coeff* and width *w* are returned.

Return af::array Result of calculated coefficients.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *widths*: Array that contains all different widths.
- *coeff*: Coefficient of interest.
- *w*: Width of interest.

af::array **energyRatioByChunks** (af::array *tss*, long *numSegments*, long *segmentFocus*)

Calculates the sum of squares of chunk i out of N chunks expressed as a ratio with the sum of squares over the whole series. *segmentFocus* should be lower than the number of segments.

Return af::array The energy ratio by chunk of the time series.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *numSegments*: The number of segments to divide the series into.
- *segmentFocus*: The segment number (starting at zero) to return a feature on.

af::array **fftAggregated** (af::array *tss*)

Calculates the spectral centroid (mean), variance, skew, and kurtosis of the absolute fourier transform spectrum.

Return af::array The spectral centroid (mean), variance, skew, and kurtosis of the absolute fourier transform spectrum.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

void **fftCoefficient** (af::array *tss*, long *coefficient*, af::array &*real*, af::array &*imag*, af::array &*abs*, af::array &*angle*)

Calculates the fourier coefficients of the one-dimensional discrete Fourier Transform for real input by using fast fourier transformation algorithm,

$$A_k = \sum_{m=0}^{n-1} a_m \exp \left\{ -2\pi i \frac{mk}{n} \right\}, \quad k = 0, \dots, n-1.$$

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `coefficient`: The coefficient to extract from the FFT.
- `real`: The real part of the coefficient.
- `imag`: The imaginary part of the coefficient.
- `abs`: The absolute value of the coefficient.
- `angle`: The angle of the coefficient.

af::array **firstLocationOfMaximum** (af::array *tss*)

Calculates the first relative location of the maximal value for each time series.

Return af::array The first relative location of the maximum value to the length of the time series, for each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **firstLocationOfMinimum** (af::array *tss*)

Calculates the first location of the minimal value of each time series. The position is calculated relatively to the length of the series.

Return af::array the first relative location of the minimal value of each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **friedrichCoefficients** (af::array *tss*, int *m*, float *r*)

Coefficients of polynomial $h(x)$, which has been fitted to the deterministic dynamics of Langevin model:

$$\dot{x}(t) = h(x(t)) + R(N)(0, 1)$$

as described by [1]. For short time series this method is highly dependent on the parameters.

[1] Friedrich et al., Physics Letters A 271, p. 217-222, Extracting model equations from experimental data, (2000).

Return `af::array` The coefficients for each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `m`: Order of polynom to fit for estimating fixed points of dynamics.
- `r`: Number of quantiles to use for averaging.

`af::array` **hasDuplicates** (`af::array tss`)

Computes if the input time series contain duplicated elements.

Return `af::array` Array containing True if the time series contains duplicated elements and false otherwise.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **hasDuplicateMax** (`af::array tss`)

Computes if the maximum within time series is duplicated.

Return `af::array` Array containing True if the maximum value of the time series is duplicated and false otherwise.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **hasDuplicateMin** (`af::array tss`)

Computes if the minimum of input time series is duplicated.

Return `af::array` Array containing True if the minimum of the time series is duplicated and false otherwise.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **indexMassQuantile** (`af::array tss`, float `q`)

Calculates the relative index i where $q\%$ of the mass of the time series within `tss` lie at the left of i . For example for $q = 50\%$ this feature calculator will return the mass center of the time series.

Return `af::array` The relative indices i where $q\%$ of the mass of the time series lie at the left of i .

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `q`: The quantile limit.

`af::array` **kurtosis** (`af::array tss`)

Returns the kurtosis of `tss` (calculated with the adjusted Fisher-Pearson standardized moment coefficient G_2).

Return `af::array` The kurtosis of `tss`.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **largeStandardDeviation** (`af::array tss`, `float r`)

Checks if the time series within `tss` have a large standard deviation.

$$\text{std}(x) > r * (\max(X) - \min(X)).$$

.

Return `af::array` Array containing True for those time series in `tss` that have a large standard deviation.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `r`: The threshold value.

`af::array` **lastLocationOfMaximum** (`af::array tss`)

Calculates the last location of the maximum value of each time series. The position is calculated relatively to the length of the series.

Return `af::array` The last relative location of the maximum value of each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **lastLocationOfMinimum** (`af::array tss`)

Calculates the last location of the minimum value of each time series. The position is calculated relatively to the length of the series.

Return `af::array` The last relative location of the minimum value of each series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **length** (`af::array tss`)

Returns the length of the input time series.

Return `af::array` The length of `tss`.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`void` **linearTrend** (`af::array tss`, `af::array &pvalue`, `af::array &rvalue`, `af::array &intercept`, `af::array &slope`, `af::array &stdr`)

Calculate a linear least-squares regression for the values of the time series versus the sequence from 0 to length of the time series minus one.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `pvalue`: The p-values for all time series.
- `rvalue`: The r-values for all time series.
- `intercept`: The intercept values for all time series.
- `slope`: The slope for all time series.
- `stderr`: The stderr values for all time series.

`af::array localMaximals (af::array tss)`

Calculates all Local Maximals for the time series in `tss`.

Return `af::array` The calculated local maxima for each time series in `tss`.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array longestStrikeAboveMean (af::array tss)`

Calculates the length of the longest consecutive subsequence in `tss` that is bigger than the mean of `tss`.

Return `af::array` the length of the longest consecutive subsequence in the input time series that is bigger than the mean.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array longestStrikeBelowMean (af::array tss)`

Calculates the length of the longest consecutive subsequence in `tss` that is below the mean of `tss`.

Return `af::array` The length of the longest consecutive subsequence in the input time series that is below the mean.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array maxLangevinFixedPoint (af::array tss, int m, float r)`

Largest fixed point of dynamics $\max_x h(x) = 0$ estimated from polynomial $h(x)$, which has been fitted to the deterministic dynamics of Langevin model:

$$\dot{x}(t) = h(x(t)) + R(N)(0, 1)$$

.

[1] Friedrich et al., Extracting model equations from experimental data, Physics Letters A 271, p. 217-222, (2000).

Return `af::array` Largest fixed point of deterministic dynamics.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series. NOTE: the time series should be sorted.
- `m`: Order of polynomial to fit for estimating fixed points of dynamics.
- `r`: Number of quantiles to use for averaging.

`af::array maximum (af::array tss)`

Calculates the maximum value for each time series within tss.

Return `af::array` The maximum value of each time series within tss.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array mean (af::array tss)`

Calculates the mean value for each time series within tss.

Return `af::array` The mean value of each time series within tss.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array meanAbsoluteChange (af::array tss)`

Calculates the mean over the absolute differences between subsequent time series values in tss.

$$\frac{1}{n} \sum_{i=1, \dots, n-1} |x_{i+1} - x_i|.$$

.

Return `af::array` The mean over the absolute differences between subsequent time series values.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array meanChange (af::array tss)`

Calculates the mean over the differences between subsequent time series values in tss.

$$\frac{1}{n} \sum_{i=1, \dots, n-1} x_{i+1} - x_i.$$

.

Return `af::array` The mean over the differences between subsequent time series values.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **meanSecondDerivativeCentral** (af::array *tss*)

Calculates mean value of a central approximation of the second derivative for each time series in *tss*.

$$\frac{1}{n} \sum_{i=1, \dots, n-1} \frac{1}{2} (x_{i+2} - 2 \cdot x_{i+1} + x_i).$$

.

Return af::array The mean value of a central approximation of the second derivative for each time series.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **median** (af::array *tss*)

Calculates the median value for each time series within *tss*.

Return af::array The median value of each time series within *tss*.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **minimum** (af::array *tss*)

Calculates the minimum value for each time series within *tss*.

Return af::array The minimum value of each time series within *tss*.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **numberCrossingM** (af::array *tss*, int *m*)

Calculates the number of *m*-crossings. A *m*-crossing is defined as two sequential values where the first value is lower than *m* and the next is greater, or viceversa. If you set *m* to zero, you will get the number of zero crossings.

Return af::array The number of *m*-crossings of each time series within *tss*.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *m*: The *m* value.

af::array **numberCwtPeaks** (af::array *tss*, int *maxW*)

This feature calculator searches for different peaks. To do so, the time series is smoothed by a ricker wavelet and for widths ranging from 1 to *maxW*. This feature calculator returns the number of peaks that occur at enough width scales and with sufficiently high Signal-to-Noise-Ratio (SNR).

Return af::array The number of peaks for each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `maxW`: The maximum width to consider.

`af::array` **numberPeaks** (`af::array tss`, `int n`)

Calculates the number of peaks of at least support n in the time series tss . A peak of support n is defined as a subsequence of tss where a value occurs, which is bigger than its n neighbours to the left and to the right.

[1] Bioinformatics (2006) 22 (17): 2059-2065. doi: 10.1093/bioinformatics/btl355, <http://bioinformatics.oxfordjournals.org/content/22/17/2059.long>

Return `af::array` The number of peaks of at least support n .

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `n`: The support of the peak.

`af::array` **partialAutocorrelation** (`af::array tss`, `af::array lags`)

Calculates the value of the partial autocorrelation function at the given lag. The lag k partial autocorrelation of a time series $\{x_t, t = 1 \dots T\}$ equals the partial correlation of x_t and x_{t-k} , adjusted for the intermediate variables $\{x_{t-1}, \dots, x_{t-k+1}\}$ ([1]). Following [2], it can be defined as:

$$\alpha_k = \frac{Cov(x_t, x_{t-k} | x_{t-1}, \dots, x_{t-k+1})}{\sqrt{Var(x_t | x_{t-1}, \dots, x_{t-k+1}) Var(x_{t-k} | x_{t-1}, \dots, x_{t-k+1})}}$$

with (a) $x_t = f(x_{t-1}, \dots, x_{t-k+1})$ and (b) $x_{t-k} = f(x_{t-1}, \dots, x_{t-k+1})$ being AR(k-1) models that can be fitted by OLS. Be aware that in (a), the regression is done on past values to predict x_t whereas in (b), future values are used to calculate the past value x_{t-k} . It is said in [1] that, for an AR(p), the partial autocorrelations α_k will be nonzero for $k \leq p$ and zero for $k > p$. With this property, it is used to determine the lag of an AR-Process.

[1] Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time series analysis: forecasting and control. John Wiley & Sons.

[2] <https://onlinecourses.science.psu.edu/stat510/node/62>

Return `af::array` The partial autocorrelation for each time series for the given lag.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `lags`: Indicates the lags to be calculated.

`af::array` **percentageOfReoccurringDatapointsToAllDatapoints** (`af::array tss`, `bool is-Sorted = false`)

Calculates the percentage of unique values, that are present in the time series more than once.

$$\frac{len(\text{different values occurring more than once})}{len(\text{different values})}$$

This means the percentage is normalized to the number of unique values, in contrast to the percentage-OfReoccurringValuesToAllValues.

Return af::array The percentage of unique data points, that are present in the time series more than once.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `isSorted`: Indicates if the input time series is sorted or not. Defaults to false.

af::array **percentageOfReoccurringValuesToAllValues** (af::array *tss*, bool *isSorted* = false)

Calculates the percentage of unique values, that are present in the time series more than once.

$$\frac{\text{number of data points occurring more than once}}{\text{number of all data points}}$$

This means the percentage is normalized to the number of unique values, in contrast to the percentage-OfReoccurringDatapointsToAllDatapoints.

Return af::array The percentage of unique values, that are present in the time series more than once.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `isSorted`: Indicates if the input time series is sorted or not. Defaults to false.

af::array **quantile** (af::array *tss*, af::array *q*, float *precision* = 100000000)

Returns values at the given quantile.

Return af::array Values at the given quantile.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `q`: Percentile(s) at which to extract score(s). One or many.
- `precision`: Number of decimals expected.

af::array **rangeCount** (af::array *tss*, float *min*, float *max*)

Counts observed values within the interval [min, max).

Return af::array Values at the given range.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `min`: Value that sets the lower limit.
- `max`: Value that sets the upper limit.

af::array **ratioBeyondRSigma** (af::array *tss*, float *r*)

Calculates the ratio of values that are more than $r * std(x)$ (so r sigma) away from the mean of x .

Return af::array The ratio of values that are more than $r * std(x)$ (so r sigma) away from the mean of x .

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `r`: Number of times that the values should be away from.

af::array **ratioValueNumberToTimeSeriesLength** (af::array *tss*)

Calculates a factor which is 1 if all values in the time series occur only once, and below one if this is not the case. In principle, it just returns:

$$\frac{\text{number_unique_values}}{\text{number_values}}$$

Return af::array The ratio of unique values with respect to the total number of values.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **sampleEntropy** (af::array *tss*)

Calculates a vectorized sample entropy algorithm. For short time-series this method is highly dependent on the parameters, but should be stable for $N > 2000$, see:

[1] Yentes et al., The Appropriate Use of Approximate Entropy and Sample Entropy with Short Data Sets, (2012).

[2] Richman & Moorman, Physiological time-series analysis using approximate entropy and sample entropy, (2000).

[3] https://en.wikipedia.org/wiki/Sample_entropy

[4] <https://www.ncbi.nlm.nih.gov/pubmed/10843903?dopt=Abstract>

Return af::array With the same dimensions as *tss*, whose values (time series in dimension 0) contains the vectorized sample entropy for all the input time series in *tss*.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **skewness** (af::array *tss*)

Calculates the sample skewness of *tss* (calculated with the adjusted Fisher-Pearson standardized moment coefficient G1).

Return af::array Containing the skewness of each time series in *tss*.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **spktWelchDensity** (af::array *tss*, int *coeff*)

Estimates the cross power spectral density of the time series *tss* at different frequencies. To do so, the time series is first shifted from the time domain to the frequency domain. Welch's method computes an estimate of the power spectral density by dividing the data into overlapping segments, computing a modified periodogram for each segment and averaging the periodograms.

[1] P. Welch, “The use of the fast Fourier transform for the estimation of power spectra: A method based on time

averaging over short, modified periodograms”, IEEE Trans. Audio Electroacoust. vol. 15, pp. 70-73, 1967.

[2] M.S. Bartlett, “Periodogram Analysis and Continuous Spectra”, Biometrika, vol. 37, pp. 1-16, 1950.

[3] Rabiner, Lawrence R., and B. Gold. “Theory and Application of Digital Signal Processing” Prentice-Hall, pp. 414-419, 1975.

Return `af::array` Containing the power spectrum of the different frequencies for each time series in `tss`.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `coeff`: The coefficient to be returned.

`af::array` **standardDeviation** (`af::array tss`)

Calculates the standard deviation of each time series within `tss`.

Return `af::array` The standard deviation of each time series within `tss`.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array` **sumOfReoccurringDatapoints** (`af::array tss`, `bool isSorted = false`)

Calculates the sum of all data points, that are present in the time series more than once.

Return `af::array` The sum of all data points, that are present in the time series more than once.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `isSorted`: Indicates if the input time series is sorted or not. Defaults to false.

`af::array` **sumOfReoccurringValues** (`af::array tss`, `bool isSorted = false`)

Calculates the sum of all values, that are present in the time series more than once.

Return `af::array` Returns the sum of all values, that are present in the time series more than once.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `isSorted`: Indicates if the input time series is sorted or not. Defaults to false.

`af::array` **sumValues** (`af::array tss`)

Calculates the sum over the time series `tss`.

Return `af::array` An array containing the sum of values in each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array symmetryLooking (af::array tss, float r)`

Calculates if the distribution of *tss* *looks symmetric*. This is the case if

$$|\text{mean}(tss) - \text{median}(tss)| < r * (\text{max}(tss) - \text{min}(tss)).$$

.

Return `af::array` Denoting if the input time series look symmetric.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `r`: The percentage of the range to compare with.

`af::array timeReversalAsymmetryStatistic (af::array tss, int lag)`

This function calculates the value of:

$$\frac{1}{n - 2lag} \sum_{i=0}^{n-2lag} x_{i+2 \cdot lag}^2 \cdot x_{i+lag} - x_{i+lag} \cdot x_i^2,$$

which is:

$$\mathbb{E}[L^2(X)^2 \cdot L(X) - L(X) \cdot X^2],$$

where \mathbb{E} is the mean and L is the lag operator. It was proposed in [1] as a promising feature to extract from time series.

[1] Fulcher, B.D., Jones, N.S. (2014). Highly comparative feature-based time-series classification. Knowledge and Data Engineering, IEEE Transactions on 26, 3026–3037.

Return `af::array` Containing the time reversal asymmetry statistic value in each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `lag`: The lag to be computed.

`af::array valueCount (af::array tss, float v)`

Counts occurrences of value in the time series *tss*.

Return `af::array` Containing the count of the given value in each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `v`: The value to be counted.

`af::array variance (af::array tss)`

Computes the variance for the time series *tss*.

Return af::array An array containing the variance in each time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **varianceLargerThanStandardDeviation** (af::array *tss*)

Calculates if the variance of *tss* is greater than the standard deviation. In other words, if the variance of *tss* is larger than 1.

Return af::array Denoting if the variance of *tss* is greater than the standard deviation.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

2.6 Namespace Library

namespace library

Typedefs

typedef *khiva_backend* Backend

Enums

enum **khiva_backend**

Values:

KHIVA_BACKEND_DEFAULT = af::Backend::AF_BACKEND_DEFAULT
Default backend order: OpenCL -> CUDA -> CPU.

KHIVA_BACKEND_CPU = af::Backend::AF_BACKEND_CPU
CPU a.k.a sequential algorithms.

KHIVA_BACKEND_CUDA = af::Backend::AF_BACKEND_CUDA
CUDA Compute Backend.

KHIVA_BACKEND_OPENCL = af::Backend::AF_BACKEND_OPENCL
OpenCL Compute Backend.

Functions

std::string **backendInfo** ()

Get information from the active backend.

Return std::string The information of the backend.

void **setBackend** (khiva::library::Backend *be*)

Set the backend.

Parameters

- `be`: The desired backend.

khiva::*library::Backend* **getBackend** ()

Get the active backend.

Return *khiva::library::Backend* The active backend.

int **getBackends** ()

Get the available backends.

Return int The available backends.

void **setDevice** (int *device*)

Set the device.

Parameters

- `device`: The desired device.

int **getDevice** ()

Get the active device.

Return int The active device.

int **getDeviceCount** ()

Get the device count.

Return int The device count.

2.7 Namespace LinAlg

namespace linalg

Functions

af::array **lls** (af::array *A*, af::array *b*)

Calculates the minimum norm least squares solution x ($\|Ax - b\|^2$) to $Ax = b$. This function uses the singular value decomposition function of Arrayfire. The actual formula that this function computes is $x = VD^\dagger U^T b$. Where U and V are orthogonal matrices and D^\dagger contains the inverse values of the singular values contained in D if they are not zero, and zero otherwise.

Return af::array Contains the solution to the linear equation problem minimizing the norm 2.

Parameters

- `A`: Coefficient matrix containing the coefficients of the linear equation problem to solve.
- `b`: Vector with the measured values.

2.8 Namespace Matrix

namespace matrix

Functions

af::array **slidingDotProduct** (af::array *q*, af::array *t*)

Calculates the sliding dot product of the time series 'q' against t.

Return array Returns an array with as many elements as 't' in the first dimension and as many elements as the last dimension of 'q' in the last dimension.

Parameters

- *q*: Array whose first dimension is the length of the query time series and the last dimension is the number of time series to calculate.
- *t*: Array with the second time series in the first dimension.

void **meanStdev** (af::array *t*, af::array &*a*, long *m*, af::array &*mean*, af::array &*stdev*)

Calculates the moving average and standard deviation of the time series 't'.

Parameters

- *t*: Input time series. Multiple time series.
- *a*: Auxiliary array to be used in the function calculateDistanceProfile. Use the overloaded method without this parameter.
- *m*: Window size.
- *mean*: Output array containing the moving average.
- *stdev*: Output array containing the moving standard deviation.

void **meanStdev** (af::array *t*, long *m*, af::array &*mean*, af::array &*stdev*)

Calculates the moving average and standard deviation of the time series 't'.

Parameters

- *t*: Input time series. Multiple time series.
- *m*: Window size.
- *mean*: Output array containing the moving average.
- *stdev*: Output array containing the moving standard deviation.

af::array **generateMask** (long *m*, long *batchSize*, long *batchStart*, long *tsLength*, long *nTimeSeries* = 1)

Function to generate a band matrix of batchSizeXtsLength with the offset batchStart.

Return af::array With the resulting band.

Parameters

- *m*: Subsequence length used to generate a band of m/2 at each side.
- *batchSize*: Size of the first dimension.

- `batchStart`: Offset of the band matrix.
- `tsLength`: Size of the second dimension of the matrix.
- `nTimeSeries`: Number of time series to generate the mask for.

void **calculateDistanceProfile** (af::array *qt*, af::array *a*, af::array *sum_q*, af::array *sum_q2*,
af::array *mean_t*, af::array *sigma_t*, af::array *mask*, af::array
&*distance*, af::array &*index*)

Calculates the distance between 'q' and the time series 't', which produced the sliding. Multiple queries can be computed simultaneously in the last dimension of 'q'.

Parameters

- `qt`: The sliding dot product of 'q' and 't'.
- `a`: Auxiliary array computed using the `meanStdev` function. This array contains a precomputed fixed value to speed up the distance calculation.
- `sum_q`: Sum of the values contained in 'q'.
- `sum_q2`: Sum of squaring the values contained in 'q'.
- `mean_t`: Moving average of 't' using a window size equal to the number of elements in 'q'.
- `sigma_t`: Moving standard deviation of 't' using a window size equal to the number of elements in 'q'.
- `mask`: Mask band matrix to filter the trivial match of a subsequence with itself.
- `distance`: Resulting minimal distance.
- `index`: Position where the minimum is occurring.

void **calculateDistanceProfile** (af::array *qt*, af::array *a*, af::array *sum_q*, af::array *sum_q2*,
af::array *mean_t*, af::array *sigma_t*, af::array &*distance*,
af::array &*index*)

Calculates the distance between 'q' and the time series 't', which produced the sliding. Multiple queries can be computed simultaneously in the last dimension of 'q'.

Parameters

- `qt`: The sliding dot product of 'q' and 't'.
- `a`: Auxiliary array computed using the `meanStdev` function. This array contains a precomputed fixed value to speed up the distance calculation.
- `sum_q`: Sum of the values contained in 'q'.
- `sum_q2`: Sum of squaring the values contained in 'q'.
- `mean_t`: Moving average of 't' using a window size equal to the number of elements in 'q'.
- `sigma_t`: Moving standard deviation of 't' using a window size equal to the number of elements in 'q'.
- `distance`: Resulting minimal distance.
- `index`: Position where the minimum is occurring.

void **mass** (af::array *q*, af::array *t*, af::array *a*, af::array *mean_t*, af::array *sigma_t*, af::array *mask*, af::array *&distance*, af::array *&index*)
Calculates the Mueen distance.

[1] Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, Hoang Anh Dau, Diego Furtado Silva, Abdullah Mueen, Eamonn Keogh (2016). Matrix Profile I: All Pairs Similarity Joins for Time Series: A Unifying View that Includes Motifs, Discords and Shapelets. IEEE ICDM 2016.

Parameters

- *q*: Array whose first dimension is the length of the query time series and the last dimension is the number of time series to calculate.
- *t*: Array with the second time series in the first dimension.
- *a*: Auxiliary array computed using the meanStdev function. This array contains a precomputed fixed value to speed up the distance calculation.
- *mean_t*: Moving average of 't' using a window size equal to the number of elements in 'q'.
- *sigma_t*: Moving standard deviation of 't' using a window size equal to the number of elements in 'q'.
- *mask*: Specifies the elements that should not be considered in the computation.
- *distance*: Resulting minimal distance.
- *index*: Position where the minimum is occurring.

void **mass** (af::array *q*, af::array *t*, af::array *a*, af::array *mean_t*, af::array *sigma_t*, af::array *&distance*, af::array *&index*)
Calculates the Mueen distance.

Parameters

- *q*: Array whose first dimension is the length of the query time series and the last dimension is the number of time series to calculate.
- *t*: Array with the second time series in the first dimension.
- *a*: Auxiliary array computed using the meanStdev function. This array contains a precomputed fixed value to speed up the distance calculation.
- *mean_t*: Moving average of 't' using a window size equal to the number of elements in 'q'.
- *sigma_t*: Moving standard deviation of 't' using a window size equal to the number of elements in 'q'.
- *distance*: Resulting minimal distance.
- *index*: Position where the minimum is occurring.

void **stomp** (af::array *ta*, af::array *tb*, long *m*, af::array *&profile*, af::array *&index*)
STOMP algorithm to calculate the matrix profile between 'ta' and 'tb' using a subsequence length of 'm'.

[1] Yan Zhu, Zachary Zimmerman, Nader Shakibay Senobari, Chin-Chia Michael Yeh, Gareth Funning, Abdullah Mueen, Philip Brisk and Eamonn Keogh (2016). Matrix Profile II: Exploiting a Novel Algorithm and GPUs to break the one Hundred Million Barrier for Time Series Motifs and Joins. IEEE ICDM 2016.

Parameters

- *ta*: Query time series.

- `tb`: Reference time series.
- `m`: Subsequence length.
- `profile`: The matrix profile, which reflects the distance to the closer element of the subsequence from 'ta' in 'tb'.
- `index`: The matrix profile index, which points to where the aforementioned minimum is located.

void **stomp** (af::array *t*, long *m*, af::array &*profile*, af::array &*index*)

STOMP algorithm to calculate the matrix profile between 't' and itself using a subsequence length of 'm'. This method filters the trivial matches.

[1] Yan Zhu, Zachary Zimmerman, Nader Shakibay Senobari, Chin-Chia Michael Yeh, Gareth Funning, Abdullah Mueen, Philip Brisk and Eamonn Keogh (2016). Matrix Profile II: Exploiting a Novel Algorithm and GPUs to break the one Hundred Million Barrier for Time Series Motifs and Joins. IEEE ICDM 2016.

Parameters

- `t`: Query and reference time series.
- `m`: Subsequence length.
- `profile`: The matrix profile, which reflects the distance to the closer element of the subsequence from 't' in a different location of itself.
- `index`: The matrix profile index, which points to where the aforementioned minimum is located.

void **findBestNMotifs** (af::array *profile*, af::array *index*, long *m*, long *n*, af::array &*motifs*, af::array &*motifsIndices*, af::array &*subsequenceIndices*, bool *selfJoin* = false)

This function extracts the best N motifs from a previously calculated matrix profile.

Parameters

- `profile`: The matrix profile containing the minimum distance of each subsequence.
- `index`: The matrix profile index containing where each minimum occurs.
- `m`: Subsequence length value used to calculate the input matrix profile.
- `n`: Number of motifs to extract.
- `motifs`: The distance of the best N motifs.
- `motifsIndices`: The indices of the best N motifs.
- `subsequenceIndices`: The indices of the query sequences that produced the minimum reported in the motifs output array.
- `selfJoin`: Indicates whether the input profile comes from a self join operation or not. It determines whether the mirror similar region is included in the output or not.

void **findBestNDiscords** (af::array *profile*, af::array *index*, long *m*, long *n*, af::array &*discords*, af::array &*discordsIndices*, af::array &*subsequenceIndices*, bool *selfJoin* = false)

This function extracts the best N discords from a previously calculated matrix profile.

Parameters

- `profile`: The matrix profile containing the minimum distance of each subsequence.
- `index`: The matrix profile index containing where each minimum occurs.

- `m`: Subsequence length value used to calculate the input matrix profile.
- `n`: Number of discords to extract.
- `discords`: The distance of the best N discords.
- `discordsIndices`: The indices of the best N discords.
- `subsequenceIndices`: The indices of the query sequences that produced the discords reported in the discords output array.
- `selfJoin`: Indicates whether the input profile comes from a self join operation or not. It determines whether the mirror similar region is included in the output or not.

2.9 Namespace Normalization

namespace normalization

Functions

`af::array decimalScalingNorm (af::array tss)`

Normalizes the given time series according to its maximum value and adjusts each value within the range (-1, 1).

Return `af::array` An array with the same dimensions as `tss`, whose values (time series in dimension 0) have been normalized by dividing each number by 10^j , where j is the number of integer digits of the max number in the time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`void decimalScalingNormInPlace (af::array &tss)`

Same as `decimalScalingNorm`, but it performs the operation in place, without allocating further memory.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

`af::array maxMinNorm (af::array tss, double high = 1.0, double low = 0.0, double epsilon = 0.00000001)`

Normalizes the given time series according to its minimum and maximum value and adjusts each value within the range [low, high].

Return `af::array` An array with the same dimensions as `tss`, whose values (time series in dimension 0) have been normalized by maximum and minimum values, and scaled as per high and low parameters.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `high`: Maximum final value (Defaults to 1.0).
- `low`: Minimum final value (Defaults to 0.0).

- `epsilon`: Safeguard for constant (or near constant) time series as the operation implies a unit scale operation between min and max values in the `tss`.

void **maxMinNormInPlace** (af::array &*tss*, double *high* = 1.0, double *low* = 0.0, double *epsilon* = 0.00000001)

Same as `maxMinNorm`, but it performs the operation in place, without allocating further memory.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `high`: Maximum final value (Defaults to 1.0).
- `low`: Minimum final value (Defaults to 0.0).
- `epsilon`: Safeguard for constant (or near constant) time series as the operation implies a unit scale operation between min and max values in the `tss`.

af::array **meanNorm** (af::array *tss*)

Normalizes the given time series according to its maximum-minimum value and its mean. It follows the following formulae:

$$\hat{x} = \frac{x - \text{mean}(x)}{\text{max}(x) - \text{min}(x)}.$$

Return af::array An array with the same dimensions as `tss`, whose values (time series in dimension 0) have been normalized by subtracting the mean from each number and dividing each number by $\text{max}(x) - \text{min}(x)$, in the time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

void **meanNormInPlace** (af::array &*tss*)

Normalizes the given time series according to its maximum-minimum value and its mean. It follows the following formulae:

$$\hat{x} = \frac{x - \text{mean}(x)}{\text{max}(x) - \text{min}(x)}.$$

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **znorm** (af::array *tss*, double *epsilon* = 0.00000001)

Calculates a new set of timeseries with zero mean and standard deviation one.

Return af::array With the same dimensions as `tss` where the time series have been adjusted for zero mean and one as standard deviation.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `epsilon`: Minimum standard deviation to consider. It acts as a gatekeeper for those time series that may be constant or near constant.

void **znormInPlace** (af::array &*tss*, double *epsilon* = 0.00000001)

Adjusts the time series in the given input and performs z-norm inplace (without allocating further memory).

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `epsilon`: Minimum standard deviation to consider. It acts as a gatekeeper for those time series that may be constant or near constant.

2.10 Namespace Polynomial

namespace polynomial

Functions

af::array **polyfit** (af::array *x*, af::array *y*, int *deg*)

Least squares polynomial fit. Fit a polynomial $p(x) = p[0] * x^{deg} + \dots + p[deg]$ of degree *deg* to points (x, y) . Returns a vector of coefficients *p* that minimizes the squared error.

Return af::array Polynomial coefficients, highest power first.

Parameters

- `x`: x-coordinates of the M sample points $(x[i], y[i])$.
- `y`: y-coordinates of the sample points.
- `deg`: Degree of the fitting polynomial.

af::array **roots** (af::array *pp*)

Calculates the roots of a polynomial with coefficients given in *p*. The values in the rank-1 array *p* are coefficients of a polynomial. If the length of *p* is $n + 1$ then the polynomial is described by:

$$p[0] * x^n + p[1] * x^{n-1} + \dots + p[n-1] * x + p[n]$$

.

Return af::array Containing the roots of the polynomial.

Parameters

- `pp`: Array of polynomial coefficients.

2.11 Namespace Regression

namespace regression

Functions

void **linear** (af::array *xss*, af::array *yss*, af::array &*slope*, af::array &*intercept*, af::array &*rvalue*, af::array &*pvalue*, af::array &*stderrest*)

Calculate a linear least-squares regression for two sets of measurements. Both arrays should have the same length.

Parameters

- *xss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *yss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *slope*: Slope of the regression line.
- *intercept*: Intercept of the regression line.
- *rvalue*: Correlation coefficient.
- *pvalue*: Two-sided p-value for a hypothesis test whose null hypothesis is that the slope is zero, using Wald Test with t-distribution of the test statistic.
- *stderrest*: Standard error of the estimated gradient.

2.12 Namespace Regularization

namespace regularization

Functions

af::array **groupBy** (af::array *in*, af::array (**aggregationFunction*)) **const** af::array&, bool, **const** dim_t, int *nColumnsKey* = 1, int *nColumnsValue* = 1 Group by operation in the input array using *nColumnsKey* columns as group keys and *nColumnsValue* columns as values. The data is expected to be sorted. The aggregation function determines the operation to aggregate the values.

Return af::array Array with the values of the group keys aggregated using the aggregationFunction.

Parameters

- *in*: Input array containing the keys and values to operate with.
- *aggregationFunction*: This param determines the operation aggregating the values.
- *nColumnsKey*: Number of columns conforming the key.
- *nColumnsValue*: Number of columns conforming the value (they are expected to be consecutive to the column keys).

af::array **groupBy** (af::array *in*, af::array (**aggregationFunction*)) **const** af::array&, **const** int, int *nColumnsKey* = 1, int *nColumnsValue* = 1 Group by operation in the input array using *nColumnsKey* columns as group keys and *nColumnsValue* columns as values. The data is expected to be sorted. The aggregation function determines the operation to aggregate the values.

Return af::array Array with the values of the group keys aggregated using the aggregationFunction.

Parameters

- `in`: Input array containing the keys and values to operate with.
- `aggregationFunction`: This param determines the operation aggregating the values.
- `nColumnsKey`: Number of columns conforming the key.
- `nColumnsValue`: Number of columns conforming the value (they are expected to be consecutive to the column keys).

`af::array groupBy (af::array in, af::array (*aggregationFunction)) const af::array&, const dim_t , int nColumnsKey = 1, int nColumnsValue = 1` Group by operation in the input array using `nColumnsKey` columns as group keys and `nColumnsValue` columns as values. The data is expected to be sorted. The aggregation function determines the operation to aggregate the values.

Return `af::array` Array with the values of the group keys aggregated using the `aggregationFunction`.

Parameters

- `in`: Input array containing the keys and values to operate with.
- `aggregationFunction`: This param determines the operation aggregating the values.
- `nColumnsKey`: Number of columns conforming the key.
- `nColumnsValue`: Number of columns conforming the value (they are expected to be consecutive to the column keys).

2.13 Namespace Statistics

`namespace statistics`

Functions

`af::array covariance (af::array tss, bool unbiased = true)`

Returns the covariance matrix of the time series contained in `tss`.

Return `af::array` The covariance matrix of the time series.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- `unbiased`: Determines whether it divides by `n - 1` (if false) or `n` (if true).

`af::array kurtosis (af::array tss)`

Returns the kurtosis of `tss` (calculated with the adjusted Fisher-Pearson standardized moment coefficient `G2`).

Return `af::array` The kurtosis of `tss`.

Parameters

- `tss`: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **moment** (af::array *tss*, int *k*)

Returns the *k*th moment of the given time series.

Return af::array The *k*th moment of the given time series.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *k*: The specific moment to be calculated.

af::array **ljungBox** (af::array *tss*, long *lags*)

The Ljung–Box test checks that data within the time series are independently distributed (i.e. the correlations in the population from which the sample is taken are 0, so that any observed correlations in the data result from randomness of the sampling process). Data are no independently distributed, if they exhibit serial correlation.

The test statistic is:

$$Q = n(n+2) \sum_{k=1}^h \frac{\hat{\rho}_k^2}{n-k}$$

where “*n*” is the sample size, $\hat{\rho}_k$ is the sample autocorrelation at lag “*k*”, and “*h*” is the number of lags being tested. Under H_0 the statistic *Q* follows a $\chi^2(h)$. For significance level α , the *criticalregion* for rejection of the hypothesis of randomness is:

$$Q > \chi_{1-\alpha, h}^2$$

where $\chi_{1-\alpha, h}^2$ is the α -quantile of the chi-squared distribution with “*h*” degrees of freedom.

[1] G. M. Ljung G. E. P. Box (1978). On a measure of lack of fit in time series models. *Biometrika*, Volume 65, Issue 2, 1 August 1978, Pages 297–303.

Return af::array Ljung-Box statistic test.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.
- *lags*: Number of lags being tested.

af::array **quantile** (af::array *tss*, af::array *q*, float *precision* = 100000000)

Returns values at the given quantile.

Return af::array Values at the given quantile.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series. NOTE: the time series should be sorted.
- *q*: Percentile(s) at which to extract score(s). One or many.
- *precision*: Number of decimals expected.

af::array **quantilesCut** (af::array *tss*, float *quantiles*, float *precision* = 0.00000001)

Discretizes the time series into equal-sized buckets based on sample quantiles.

Return af::array Matrix with the categories, one category per row, the start of the category in the first column and the end in the second category.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series. NOTE: the time series should be sorted.
- *quantiles*: Number of quantiles to extract. From 0 to 1, step 1/quantiles.
- *precision*: Number of decimals expected.

af::array **sampleStdev** (af::array *tss*)

Estimates standard deviation based on a sample. The standard deviation is calculated using the “n-1” method.

Return af::array The sample standard deviation.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

af::array **skewness** (af::array *tss*)

Calculates the sample skewness of *tss* (calculated with the adjusted Fisher-Pearson standardized moment coefficient G1).

Return af::array Array containing the skewness of each time series in *tss*.

Parameters

- *tss*: Expects an input array whose dimension zero is the length of the time series (all the same) and dimension one indicates the number of time series.

- *Namespace Array*
- *Namespace Clustering*
- *Namespace Dimensionality*
- *Namespace Distances*
- *Namespace Features*
- *Namespace Library*
- *Namespace LinAlg*
- *Namespace Matrix*
- *Namespace Normalization*
- *Namespace Polynomial*
- *Namespace Regression*
- *Namespace Regularization*
- *Namespace Statistics*

We have developed bindings to enable the execution of Khiva from the following languages. In order to make it work, you should first install Khiva library in your machine, explained in :ref: *chapter-gettingstarted*.

3.1 Python

In order to install the khiva-python binding of the library, you would need to fetch the latest version of the code from:

```
git clone https://github.com/shapelets/khiva-python.git
```

After cloning the repository, you can install khiva-python by executing the next commands:

```
cd /path_to_khiva-python
python3 setup.py install
```

If the installation is successful, you are ready to start playing with the library.

3.2 Java

In order to install the khiva-java binding of the library, you would need to fetch the latest version of the code from:

```
git clone https://github.com/shapelets/khiva-java.git
```

Once you have downloaded the code, you have to move to the source code directory and execute the following commands:

```
cd path_to_java_khiva_dir
mvn install
mvn javadoc:javadoc
```

If all steps finished as expected, you should be able to use the Khiva from your java projects.

3.3 R

In order to install the khiva-r binding of the library, you would need to fetch the latest version of the code from:

```
git clone https://github.com/shapelets/khiva-r.git
```

After downloading the code, you would need to open an R console and execute the following commands, to set the work directory and install the Khiva binding:

```
setwd(<project-root-dir>/)
devtools::install()
```

Once the installation of the binding has been carried out, you can make the library available by executing:

```
library(khiva)
```

If all previous steps were successful you will be ready to start working with the library.

3.4 MATLAB

In order to install the khiva-matlab binding of the library, you would need to fetch the latest version of the code from:

```
git clone https://github.com/shapelets/khiva-matlab.git
```

Once the code is available, we just have to add the path to the khiva-matlab/+khiva folder to the MATLAB path. Thus, the user will be able to import and call our library.

Building programs using Khiva with CMake

In order to build a program using the Khiva library with the [CMake](#) build system you need just a couple of lines in your `CmakeLists.txt`:

```
cmake_minimum_required(VERSION 3.1)
project(example)

find_package(Khiva REQUIRED)

add_executable(example example.cpp)
target_link_libraries(example Khiva::khiva)
```

`find_package(Khiva REQUIRED)` may be used when Khiva was installed system wide. Please follow the installation instructions for your operating system contained at the [Getting Started](#).

5.1 Core Development Team

- Justo Ruiz-Ferrer (justo.ruiz@shapelets.io)
- Antonio Vilches (antonio.vilches@shapelets.io)
- Oscar Torreno (oscar.torreno@shapelets.io)
- David Cuesta (david.cuesta@shapelets.io)

5.2 Contributions

- Luis Sanchez (luis.sanchez@shapelets.io)

CHAPTER 6

Cite Us

If you use Khiva Library for a publication, please cite it as:

```
@misc{khiva-library,  
  author = "David Cuesta and Justo Ruiz and Oscar Torreno and Antonio Vilches",  
  title = "Khiva Library",  
  howpublished = "\url{https://shapelets.io/khiva}"  
}
```


K

- khiva::array (C++ type), 9
- khiva::array::Array (C++ class), 10
- khiva::array::Array::~Array (C++ function), 11
- khiva::array::Array::Array (C++ function), 11
- khiva::array::Array::getColumn (C++ function), 12
- khiva::array::Array::getData (C++ function), 12
- khiva::array::Array::getElement (C++ function), 12
- khiva::array::Array::getNumElements (C++ function), 12
- khiva::array::Array::getNumW (C++ function), 11
- khiva::array::Array::getNumX (C++ function), 11
- khiva::array::Array::getNumY (C++ function), 11
- khiva::array::Array::getNumZ (C++ function), 11
- khiva::array::Array::getRow (C++ function), 12
- khiva::array::Array::isEmpty (C++ function), 12
- khiva::array::Array::print (C++ function), 12
- khiva::array::Array::setData (C++ function), 11
- khiva::array::Array::setNumW (C++ function), 11
- khiva::array::Array::setNumX (C++ function), 11
- khiva::array::Array::setNumY (C++ function), 11
- khiva::array::Array::setNumZ (C++ function), 11
- khiva::array::createArray (C++ function), 9
- khiva::array::deleteArray (C++ function), 10
- khiva::array::getData (C++ function), 9
- khiva::array::getDims (C++ function), 9
- khiva::array::getIndexMaxColumns (C++ function), 10
- khiva::array::getRowsWithMaximals (C++ function), 10
- khiva::array::getType (C++ function), 10
- khiva::array::print (C++ function), 10
- khiva::clustering (C++ type), 12
- khiva::clustering::kMeans (C++ function), 12
- khiva::clustering::kShape (C++ function), 13
- khiva::dimensionality (C++ type), 13
- khiva::dimensionality::PAA (C++ function), 13, 14
- khiva::dimensionality::PIP (C++ function), 14
- khiva::dimensionality::PLABottomUp (C++ function), 14
- khiva::dimensionality::PLASlidingWindow (C++ function), 15
- khiva::dimensionality::Point (C++ type), 13
- khiva::dimensionality::ramerDouglasPeucker (C++ function), 15
- khiva::dimensionality::SAX (C++ function), 16
- khiva::dimensionality::Segment (C++ type), 13
- khiva::dimensionality::visvalingam (C++ function), 16
- khiva::distances (C++ type), 17
- khiva::distances::dtw (C++ function), 17
- khiva::distances::euclidean (C++ function), 17
- khiva::distances::hamming (C++ function), 17
- khiva::distances::manhattan (C++ function), 18
- khiva::distances::sbd (C++ function), 18
- khiva::distances::squaredEuclidean (C++ function), 18
- khiva::features (C++ type), 18
- khiva::features::absEnergy (C++ function), 19
- khiva::features::absoluteSumOfChanges (C++ function), 19
- khiva::features::aggregatedAutocorrelation (C++ function), 19, 20
- khiva::features::aggregatedLinearTrend (C++ function), 20, 21
- khiva::features::approximateEntropy (C++ function), 21
- khiva::features::autoCorrelation (C++ function), 21
- khiva::features::autoCovariance (C++ function), 22
- khiva::features::binnedEntropy (C++ function), 22
- khiva::features::c3 (C++ function), 22
- khiva::features::cidCe (C++ function), 23
- khiva::features::countAboveMean (C++ function), 23
- khiva::features::countBelowMean (C++ function), 23
- khiva::features::crossCorrelation (C++ function), 23
- khiva::features::crossCovariance (C++ function), 23
- khiva::features::cwtCoefficients (C++ function), 24
- khiva::features::energyRatioByChunks (C++ function), 24
- khiva::features::fftAggregated (C++ function), 24
- khiva::features::fftCoefficient (C++ function), 25
- khiva::features::firstLocationOfMaximum (C++ function), 25
- khiva::features::firstLocationOfMinimum (C++ function), 25
- khiva::features::friedrichCoefficients (C++ function), 25

khiva::features::hasDuplicateMax (C++ function), 26
khiva::features::hasDuplicateMin (C++ function), 26
khiva::features::hasDuplicates (C++ function), 26
khiva::features::indexMassQuantile (C++ function), 26
khiva::features::kurtosis (C++ function), 26
khiva::features::largeStandardDeviation (C++ function), 27
khiva::features::lastLocationOfMaximum (C++ function), 27
khiva::features::lastLocationOfMinimum (C++ function), 27
khiva::features::length (C++ function), 27
khiva::features::linearTrend (C++ function), 27
khiva::features::localMaximals (C++ function), 28
khiva::features::longestStrikeAboveMean (C++ function), 28
khiva::features::longestStrikeBelowMean (C++ function), 28
khiva::features::maximum (C++ function), 29
khiva::features::maxLangevinFixedPoint (C++ function), 28
khiva::features::mean (C++ function), 29
khiva::features::meanAbsoluteChange (C++ function), 29
khiva::features::meanChange (C++ function), 29
khiva::features::meanSecondDerivativeCentral (C++ function), 29
khiva::features::median (C++ function), 30
khiva::features::minimum (C++ function), 30
khiva::features::numberCrossingM (C++ function), 30
khiva::features::numberCwtPeaks (C++ function), 30
khiva::features::numberPeaks (C++ function), 31
khiva::features::partialAutocorrelation (C++ function), 31
khiva::features::percentageOfReoccurringDatapointsToAllDatapoints (C++ function), 31
khiva::features::percentageOfReoccurringValuesToAllValues (C++ function), 32
khiva::features::quantile (C++ function), 32
khiva::features::rangeCount (C++ function), 32
khiva::features::ratioBeyondRSigma (C++ function), 32
khiva::features::ratioValueNumberToTimeSeriesLength (C++ function), 33
khiva::features::sampleEntropy (C++ function), 33
khiva::features::skewness (C++ function), 33
khiva::features::spktWelchDensity (C++ function), 33
khiva::features::standardDeviation (C++ function), 34
khiva::features::sumOfReoccurringDatapoints (C++ function), 34
khiva::features::sumOfReoccurringValues (C++ function), 34
khiva::features::sumValues (C++ function), 34
khiva::features::symmetryLooking (C++ function), 35
khiva::features::timeReversalAsymmetryStatistic (C++ function), 35
khiva::features::valueCount (C++ function), 35
khiva::features::variance (C++ function), 35
khiva::features::varianceLargerThanStandardDeviation (C++ function), 36
khiva::library (C++ type), 36
khiva::library::Backend (C++ type), 36
khiva::library::backendInfo (C++ function), 36
khiva::library::getBackend (C++ function), 37
khiva::library::getBackends (C++ function), 37
khiva::library::getDevice (C++ function), 37
khiva::library::getDeviceCount (C++ function), 37
khiva::library::khiva_backend (C++ type), 36
khiva::library::KHIVA_BACKEND_CPU (C++ enumerator), 36
khiva::library::KHIVA_BACKEND_CUDA (C++ enumerator), 36
khiva::library::KHIVA_BACKEND_DEFAULT (C++ enumerator), 36
khiva::library::KHIVA_BACKEND_OPENCL (C++ enumerator), 36
khiva::library::setBackend (C++ function), 36
khiva::library::setDevice (C++ function), 37
khiva::linalg (C++ type), 37
khiva::linalg::lls (C++ function), 37
khiva::matrix (C++ type), 38
khiva::matrix::calculateDistanceProfile (C++ function), 39
khiva::matrix::findBestNDiscords (C++ function), 41
khiva::matrix::findBestNMotifs (C++ function), 41
khiva::matrix::generateMask (C++ function), 38
khiva::matrix::mass (C++ function), 39, 40
khiva::matrix::meanStdev (C++ function), 38
khiva::matrix::slidingDotProduct (C++ function), 38
khiva::matrix::stomp (C++ function), 40, 41
khiva::normalization (C++ type), 42
khiva::normalization::decimalScalingNorm (C++ function), 42
khiva::normalization::decimalScalingNormInPlace (C++ function), 42
khiva::normalization::maxMinNorm (C++ function), 42
khiva::normalization::maxMinNormInPlace (C++ function), 43
khiva::normalization::meanNorm (C++ function), 43
khiva::normalization::meanNormInPlace (C++ function), 43
khiva::normalization::znorm (C++ function), 43
khiva::normalization::znormInPlace (C++ function), 44
khiva::polynomial (C++ type), 44
khiva::polynomial::polyfit (C++ function), 44
khiva::polynomial::roots (C++ function), 44
khiva::regression (C++ type), 44
khiva::regression::linear (C++ function), 45
khiva::regularization (C++ type), 45
khiva::regularization::groupBy (C++ function), 45, 46
khiva::statistics (C++ type), 46

khiva::statistics::covariance (C++ function), [46](#)
khiva::statistics::kurtosis (C++ function), [46](#)
khiva::statistics::ljungBox (C++ function), [47](#)
khiva::statistics::moment (C++ function), [46](#)
khiva::statistics::quantile (C++ function), [47](#)
khiva::statistics::quantilesCut (C++ function), [48](#)
khiva::statistics::sampleStdev (C++ function), [48](#)
khiva::statistics::skewness (C++ function), [48](#)