

delivering benefits through evidence



Development of interim national
guidance on non-stationary fluvial flood
frequency estimation –
R package `nonstat` user guide

FRS18087/IG/R3

We are the Environment Agency. We protect and improve the environment.

Acting to reduce the impacts of a changing climate on people and wildlife is at the heart of everything we do.

We reduce the risks to people, properties and businesses from flooding and coastal erosion.

We protect and improve the quality of water, making sure there is enough for people, businesses, agriculture and the environment. Our work helps to ensure people can enjoy the water environment through angling and navigation.

We look after land quality, promote sustainable land management and help protect and enhance wildlife habitats. And we work closely with businesses to help them comply with environmental regulations.

We can't do this alone. We work with government, local councils, businesses, civil society groups and communities to make our environment a better place for people and wildlife.

Published by:

Environment Agency, Horizon House, Deanery Road,
Bristol, BS1 5AH

[http://www.gov.uk/government/organisations/
environment-agency](http://www.gov.uk/government/organisations/environment-agency)

ISBN: 978-1-84911-469-1

© Environment Agency – November 2020

All rights reserved. This document may be reproduced
with prior permission of the Environment Agency.

Email: fcem.evidence@environment-agency.gov.uk

Further copies of this report are available from our
publications catalogue:

<http://www.gov.uk/government/publications>

or our National Customer Contact Centre:
T: 03708 506506

Email: enquiries@environment-agency.gov.uk

Author(s):

Sarah Warren and Dr Sean Longfield

Dissemination status:

Publicly available

Keywords:

Stationarity, non-stationarity, flood frequency
estimation, trend

Research contractor:

JBA Consulting
1 Broughton Park, Old Lane North, Broughton,
Skipton, North Yorkshire, BD23 3FD

Environment Agency's Project Manager:

Dr Sean Longfield, Lead Scientist

Theme Manager:

Dr Sue Manson, Incident Management and Modelling
Theme Manager

Collaborator(s):

Dr Janet Heffernan

Evidence at the Environment Agency

Scientific research and analysis underpins everything the Environment Agency does. It helps us to understand and manage the environment effectively. Our own experts work with leading scientific organisations, universities and other parts of the Defra group to bring the best knowledge to bear on the environmental problems that we face now and in the future. Our scientific work is published as summaries and reports, freely available to all.

This report is the result of research commissioned and funded by the Joint Flood and Coastal Erosion Risk Management Research and Development Programme. The Joint Programme is jointly overseen by Defra, the Environment Agency, Natural Resources Wales and the Welsh Government on behalf of all risk management authorities in England and Wales:

<https://www.gov.uk/government/organisations/flood-and-coastal-erosion-risk-management-research-and-development-programme>

You can find out more about our current science programmes at:

<https://www.gov.uk/government/organisations/environment-agency/about/research>

If you have any comments or questions about this report or the Environment Agency's other scientific work, please contact research@environment-agency.gov.uk

Professor Doug Wilson
Director, Research, Analysis and Evaluation

Executive summary

Reports of ‘record-breaking’ or ‘unprecedented’ floods rarely seem out of the news in England and Wales. This has raised questions about whether the probability of floods has changed over the 40 to 60 years for which river flow data are typically available, and whether probability might change in the future.

Flood frequency analysis tells us what flood flows are expected to occur with a given probability. It is a fundamental part of cost-benefit analysis which is used to make decisions on investment in flood protection. Flood frequency analysis is also important for other areas of flood risk management such as mapping of flood risk for planning, long-term investment planning, national flood risk assessment, setting insurance premiums, designing river structures and for reservoir safety.

In UK flood frequency analysis practitioners use the methods in the Flood Estimation Handbook. These techniques assume that in a data series each value, for example, each annual maximum flow or rainfall, is independent and has the same probability distribution as all the other values. If this probability distribution is not constant over time, (non-stationary), the peak flows are not identically distributed and so this assumption cannot be made. This could have significant implications for capital investment decisions now and in the future if it means that flood risk has been over or underestimated.

The main aim of this project was to develop interim guidance for dealing with non-stationarity in annual maximum river flow series. The objectives were to:

- develop methods for identifying non-stationarity in annual maximum flow series
- develop a scientifically robust process for carrying out non-stationary flood frequency analysis
- outline how to take account of future climate change under non-stationary conditions
- provide guidance and tools for practitioners to carry out non-stationary flood frequency analysis
- make a high-level assessment of the impact of allowing for non-stationarity on flood frequency estimation across England and Wales

This report focuses on providing guidance and tools for practitioners to carry out non-stationary flood frequency analysis. It describes a number of functions that have been developed as part of a package in ‘R’ called `nonstat`. R is a programming language and software environment for statistical computing and graphics. An R package is a collection of code, documentation and data that allow users to easily carry out complex statistical analysis.

Three functions to detect trend and changepoints in records of annual maximum river flow are described in this report:

- `MK.test`
This function carries out the Mann-Kendall trend test on data from river flow gauges. The Mann-Kendall test is a non-parametric trend test, widely used for monotonic trend testing.
- `Pettitt.test`
This function carries out Pettitt’s test, which is designed to detect a sudden

change in the mean of a time series. It is a non-parametric test, that is, it makes no assumption about the distribution followed by the data.

- `PELT.test`
the PELT, or Pruned Exact Linear Time test is an alternative to the Pettitt test for detecting change points and allows for detecting one or more change(s) in mean and variance of annual maximum river flow data.

Five functions to carry out non-stationary flood frequency analysis, and present the results are also outlined:

- `fit.time`
This function fits stationary and non-stationary distributions (generalised extreme value (GEV) or generalised logistic (GLO)) to annual maximum river flow data. Three non-stationary models can be fitted, by varying the location parameter, varying scale parameter or varying both location and scale parameters. The parameters vary with a covariate, in this case, time.
- `plot.selected.distribution`
This function plots the results from the fitted stationary distribution and one of the fitted non-stationary distributions selected by the user (for example, the best fitting non-stationary model) for one user-selected gauge using the outputs from `fit.time` for a range of return periods. It also plots the confidence interval for one return period on a second plot.
- `plot.all`
This function plots the results from the fitted stationary and 3 non-stationary models (varying location, varying scale and varying location and scale parameters) with time as a covariate for a user-specified gauge.
- `fit.phys.cov`
This function fits stationary and non-stationary GEV or GLO distributions to annual maximum flow data. The nonstationary models have varying location and/or scale parameters that can vary with one or more physical covariate(s) and/or time. Examples of physical covariates that could be used include annual or seasonal rainfall over the catchment, temperature, large-scale indices of atmospheric circulation, such as the North Atlantic Oscillation (NAO) or East Atlantic pattern (EA) and urban extent for the catchment.
- `res.phys.cov`
This function uses one of the fitted non-stationary models from the `fit.phys.cov` function to estimate conditional, marginal and present-day marginal return levels. Stationary return levels are also estimated based on the fitted stationary model. The function also outputs plots and tables of the results.

Acknowledgements

Particular thanks are extended to the Environment Agency's project team: Richard Macilwaine, Peter Spencer, Dr Sean Longfield and Parveen Mann.

We are particularly grateful to Dr Janet Heffernan for her work on implementing the generalised logistic distribution in the `textmex` R package.

We are also grateful to Martin Borthwick, David Lindsay, Ed Malone, Peter Spencer, Mike Vaughan (of the Environment Agency), Glenda Tudor-Ward (Natural Resources Wales), Kim Hearn (Aecom), Luke Ballantyne (Arup), Phil Raynor (Jacobs), and Tracey Ashworth (Atkins) for providing constructive feedback during the testing phase of this package development.

Contents

1	Introduction	1
1.1	Conventions used in this report	1
1.2	RStudio	2
1.3	Package installation	3
2	Using the <code>nonstat</code> package	5
2.1	Data	5
2.2	General comments	6
3	Trend and change point tests	7
3.1	<code>MK.test</code>	7
3.2	<code>Pettitt.test</code>	10
3.3	<code>PELT.test</code>	12
4	Fitting non-stationary models	15
4.1	Fitting models where time is the only covariate	17
4.2	Fitting models with physical covariates	28
5	Complete example script to run all functions	39
	References	41

List of figures

Figure 1-1:	RStudio script example	2
Figure 1-2:	Example of the successful installation of <code>nonstat</code> in RStudio	4
Figure 3-1:	Example of the graphical output of the <code>MK.test</code> function	9
Figure 3-2:	Example of the graphical output of the <code>Pettitt.test</code> function	11
Figure 3-3:	Example of normal Q-Q plots from the <code>PELT.test</code> function	13
Figure 3-4:	Example of a change point from the <code>PELT.test</code> function	14
Figure 4-1:	Example diagnostic plot from the <code>fit.time</code> function: Observed AMAX data, probability-probability (or P-P) plot and a quantile-quantile (or Q-Q) plot.	21
Figure 4-2:	Example of Q-Q plots for all 4 models (time as covariate) from the <code>fit.time</code> function	21
Figure 4-3:	Example of output plots from the <code>plot.selected.distributions</code> function	23
Figure 4-4:	Example of output plots from the <code>plot.all</code> function	25
Figure 4-5:	Example of output plot of flow data against all covariates in the input file from the <code>res.phys.cov</code> function	35
Figure 4-6:	Example of output plot of flow against conditional return levels from the non-stationary model and stationary return levels from the stationary model for the observed period from the <code>res.phys.cov</code> function	36
Figure 4-7:	Example of output plot of flow against encounter probability from the <code>res.phys.cov</code> function	37

1 Introduction

The `nonstat` R package is designed to allow hydrologists to easily carry out trend and change point detection and non-stationary flood frequency analysis using annual maximum (AMAX) flow data. The package contains a number of functions that the user can call, all of which have associated help files. This document provides more detail on the functions, their inputs and outputs, and how to use them. This user guide is supported by 2 other documents. The first is practitioner guidance (FRS18087 – Allowing for non-stationarity in flood frequency estimation – Non-stationarity Science Report which aims to introduce practitioners to the concepts of non-stationary flood frequency analysis and guide them as to when and how they should do it and how to interpret the results. It illustrates the concepts with 3 case studies. The second is a science report (FRS18087 – Allowing for non-stationarity in flood frequency estimation – Non-stationarity Science Report) which gives more scientific background on the methods implemented in the `nonstat` package, as well as nationwide results from applying the methods to the gauging station network in England and Wales.

It is recommended that analysis is carried out using **RStudio**, a free and open-source integrated development environment (IDE) for R. A version of R of at least 3.5.0 is required.

This documentation contains examples of code for running each function independently and there is therefore some repetition throughout the examples of assigning the same values to the same variables. However, there is also a complete example script at the end of the document for all aspects of analysis that can be carried out using this package and this avoids the repetition. All examples here use a folder called 'Non-stationarity' saved on the C drive. However, you can change the file path to anything you wish. The analysis does not need to be carried out in the same folder in which the `nonstat` package is saved.

1.1 Conventions used in this report

Code snippets in this report appear like the example below:

```
> folder <- "C:/Non-stationarity/Data"
> Pettitt.test(folder, sig.level = 5)

# This is an example of a comment in the code
```

Each line of R code in the examples given is preceded by one of 2 symbols:

>: the prompt symbol, > is not part of the code, it is the command line from the R console, and you should not type this when you use the code in RStudio

#: the hash symbol, # is not part of the code, it is used to put comments in the code to help explain code to users

The `Courier New` font is used to highlight package names, function names, function arguments and lines of code in this report.

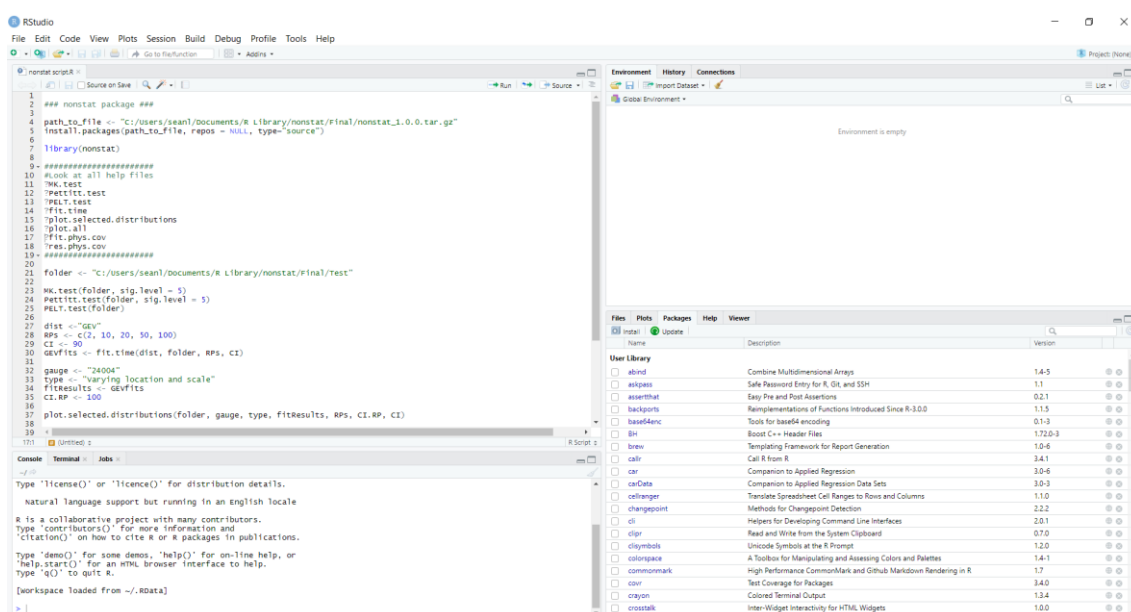
1.2 RStudio

It is recommended that analysis is carried out using RStudio. For users with admin rights to install software on their computers, RStudio can be downloaded and installed from <https://rstudio.com/products/rstudio/download/>

Environment Agency users can access RStudio via the Central Modelling Platform (CMP). R and R-Studio are part of the standard suite of application on the CMP production environment and are available to all users.

To ensure that analyses are reproducible as far as possible, it is recommended that you create and save an R script (in RStudio: File -> New File -> R Script). This script will contain the lines of code that you wish to run. The script will open in the top left quadrant (shown in Figure 1-1). The hash symbol (#) can be used to comment out parts of the script that should not be run, for example, if you want to add a comment about what you are doing, or don't want to run a certain line of code.

Figure 1-1: RStudio script example



Lines of code in the script can be run in the console by clicking on the line to run and pressing **Ctrl + Enter** or by clicking the **Run** button at the top of the script. You can save the script as a **.R** file using **File -> Save As**. This can then be re-opened if you need to access the script again.

It is generally recommended that you start a new analysis in a clean version of RStudio, where nothing is saved in the workspace that may interfere with the analysis. If you get errors that you have not previously encountered from the same function, then it is worth restarting RStudio to check that nothing is saved from previous analyses that may be interfering with the new analysis.

1.3 Package installation

The `nonstat` R package is supplied as a `.tar.gz` file.

The file name is: **`nonstat_1.0.0.tar.gz`**

The `nonstat` package has a number of dependencies (other R packages on which the code depends). These are:

- `changepoint`
- `data.table`
- `dplyr`
- `ggplot2`
- `grid`
- `gridExtra`
- `lubridate`
- `reshape2`
- `stats`
- `stringr`
- `texmex`
- `trend.`

These must be installed first. You can see which packages you've installed in the 'Packages' tab, in the bottom right panel (assuming the default RStudio layout).

Packages only need installing once, so you don't need to reinstall packages that you've already installed (unless they need updating). If you need to install any dependencies, you can do this by using the following code (listing all missing packages, rather than just the 2 in this example):

```
> install.packages(c('changepoint', 'trend'))
```

Save the `.tar.gz` file to a directory with an appropriate name. You can use the code below to store that directory and file name in R, replacing the file path with the one to the directory you just created. Note that file and folder paths in R use a forward slash rather than a backslash by default.

```
> path_to_file <- "C:/Non-stationarity/nonstat_1.0.0.tar.gz"
```

Now you can install the `nonstat` package, using the following code. (It would also be possible to put the file path directly into the below line, rather than using 'path_to_file', if preferred.)

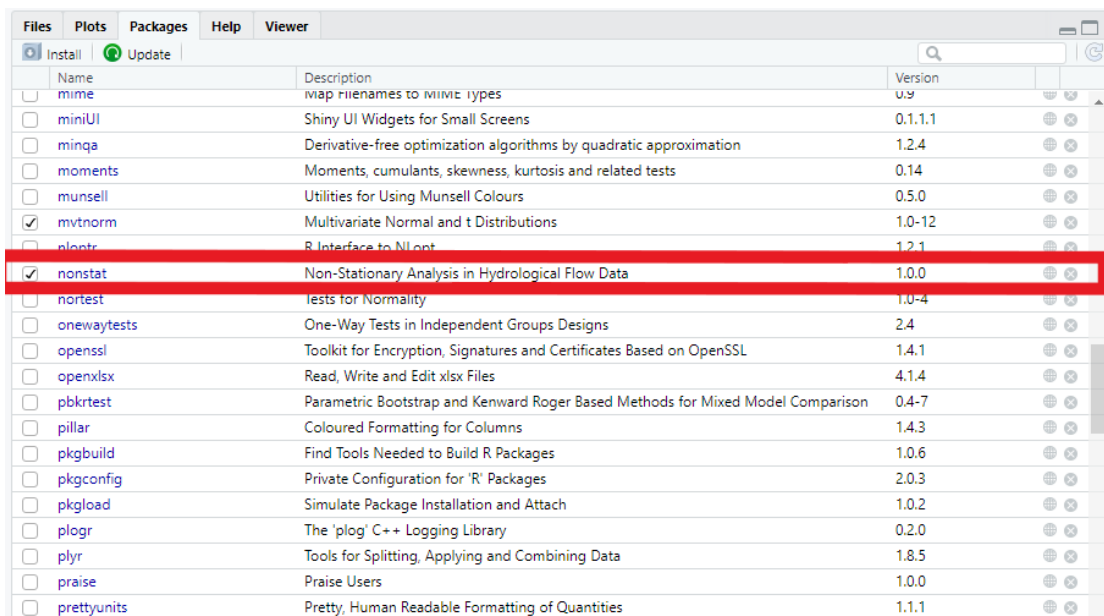
```
> install.packages(path_to_file, repos = NULL, type="source")
```

If any of the package dependencies are still missing, you will get an error message like the following. In this example, the only missing package is `changepoint`. This means that the `nonstat` package has failed to install and you'll need to install any missing packages before reinstalling the `nonstat` package.

```
ERROR: dependency 'changepoint' is not available for package 'nonstat'
* removing 'C:/Program Files/R/R-3.6.3/library/nonstat'
warning in install.packages :
  installation of package 'C:/Non-Stationary/nonstat_1.0.0.tar.gz' had
  non-zero exit status
```

Once the package is installed, you should be able to see the package present in the 'Packages' tab, in the bottom right panel (assuming the default RStudio layout), as shown below in Figure 1-2:

Figure 1-2: Example of the successful installation of `nonstat` in RStudio



Files	Plots	Packages	Help	Viewer
Install Update				
Name	Description	Version		
<input type="checkbox"/> mime	map filenames to mime types	0.9		
<input type="checkbox"/> miniUI	Shiny UI Widgets for Small Screens	0.1.1.1		
<input type="checkbox"/> minqa	Derivative-free optimization algorithms by quadratic approximation	1.2.4		
<input type="checkbox"/> moments	Moments, cumulants, skewness, kurtosis and related tests	0.14		
<input type="checkbox"/> munsell	Utilities for Using Munsell Colours	0.5.0		
<input checked="" type="checkbox"/> mvtnorm	Multivariate Normal and t Distributions	1.0-12		
<input type="checkbox"/> nlortr	R Interface to Nlort	1.2.1		
<input checked="" type="checkbox"/> nonstat	Non-Stationary Analysis in Hydrological Flow Data	1.0.0		
<input type="checkbox"/> nortest	Tests for Normality	1.0-4		
<input type="checkbox"/> onewaytests	One-Way Tests in Independent Groups Designs	2.4		
<input type="checkbox"/> openssl	Toolkit for Encryption, Signatures and Certificates Based on OpenSSL	1.4.1		
<input type="checkbox"/> openxlsx	Read, Write and Edit xlsx Files	4.1.4		
<input type="checkbox"/> pbkrtest	Parametric Bootstrap and Kenward Roger Based Methods for Mixed Model Comparison	0.4-7		
<input type="checkbox"/> pillar	Coloured Formatting for Columns	1.4.3		
<input type="checkbox"/> pkgbuild	Find Tools Needed to Build R Packages	1.0.6		
<input type="checkbox"/> pkgconfig	Private Configuration for 'R' Packages	2.0.3		
<input type="checkbox"/> pkgload	Simulate Package Installation and Attach	1.0.2		
<input type="checkbox"/> plogr	The 'plog' C++ Logging Library	0.2.0		
<input type="checkbox"/> plyr	Tools for Splitting, Applying and Combining Data	1.8.5		
<input type="checkbox"/> praise	Praise Users	1.0.0		
<input type="checkbox"/> prettyunits	Pretty, Human Readable Formatting of Quantities	1.1.1		

2 Using the `nonstat` package

You can access the functions in the `nonstat` package by loading the package into the current R session by running the following code in the console (or by checking the box to the left of the package name in the 'Packages' list shown above). You will be able to see that other R packages on which `nonstat` depends are also loaded.

```
> library(nonstat)
```

You can access the help files for each function in R by typing `?function` where function is the function name, for example:

```
> ?fit.time
```

Note that R is case-sensitive.

2.1 Data

Each function in the package is developed to work with flow data in .AM files, such as those found in the National River Flow Archive (NRFA) Peak Flow Dataset. The dates must be in the standard format of 22 Jan 1960 for example; please make sure that this is the case if you manually edit any files. Before running any functions, you need to create a directory containing 2 folders: '**Inputs**' and '**Outputs**'. To tell R what this directory is, you can use the following code:

```
> folder <- "C:/Non-stationarity/Data"
```

Every function in the package has the path to this directory as an input argument. If you are using the same folder for the whole analysis, this can be set just once at the beginning of the script.

Data for any gauges you wish to analyse must be in the 'Inputs' folder (many of the functions will loop through all the gauge data (.AM files) saved in this folder). Outputs will automatically be written out to the 'Outputs' folder.

2.2 General comments

For outputs from all functions, the gauge numbers come from the station number within the .AM file and therefore you'll need to change that if you use multiple versions of the same gauge (if only the file name of the .AM file is changed, some of the results files will be overwritten). Outputs will be overwritten when the code is re-run if the gauge numbers are the same. However, output file names can be changed once the file has been produced if needed. (Note that a .png file will be overwritten even if the file is open, but a .csv file will not. A warning message will be produced, so it's important to make sure that .csv results files are closed before running code to update them.) The gauge name comes from a look-up file ('Gauge_metadata.csv'), which is provided with the package and needs to be placed in the 'Inputs' folder. If the gauge ID (number) does not exist in the look-up file, no gauge name will be included in the outputs. The file can be edited to include your gauge if you wish to do so, as long as the format of the file remains the same. The existing IDs are those of the NRFA.

When running the scripts, it is worth keeping an eye out for errors or warnings that may appear in the R console. Errors generally stop the function so that it fails to complete its intended analysis, so it is unlikely that some or all results will have been written out. The problem will therefore need addressing. Some of the functions catch the errors, rather than stopping the analysis, and they are written out in .csv files instead. It is therefore important to check the warning and error files that are written out.

Warnings caution users without stopping the function working, so you'll probably get results, but it's possible that there was a problem with the inputs or the computation, for example, something may have been amended slightly to allow the function to work. Some warnings don't require you to do anything, if you're happy that the analysis is not being impacted and the results are correct.

Each function in this package uses the same data pre-processing routine. If the data pre-processing produces warnings or errors for a given gauge, they will be written out in the automatically-generated 'Input data warnings and errors.csv' file. Errors are likely to be related to formatting of the input data, for example, an expected header being missing or one section of the file not ending with an '[END]' row.

A warning message may also be produced if one of the gauges has a missing stage value in the .AM file. This is not a problem as it's only the flow values that are analysed, but it's worth checking that this is the problem in the .AM file to which the warning refers. In the unlikely situation that a flow value is missing in the .AM file for a date (water year) that is included in the file, that row (water year) is automatically removed from the analysis (to enable the distribution fitting in some of the functions to work).

All water years in the output files are defined as starting in the October of the year stated, for example, the 2019 water year would be from October 2019 to September 2020.

The x and y axes of all plots are determined automatically based on the data; some are fixed such that multiple plots showing results for different return periods can be easily compared.

Optional parameters do not need to be included to run the function. Output folders and files will automatically be created for each function that is run.

3 Trend and change point tests

There are 3 functions in the package that can carry out trend and change point tests and these are described in this section. All these functions can be run on multiple gauges and a list of the gauges will be printed in the R console. A progress bar will update after each gauge has been analysed.

3.1 MK.test

This function carries out the Mann-Kendall trend test on all gauges listed in a folder and uses functionality from the `trend` R package. The Mann-Kendall test is a non-parametric trend test, widely used for monotonic trend testing. The results are automatically written out to the 'Outputs/Mann-Kendall' folder. The 'Mann-Kendall' folder will be automatically generated if it does not already exist.

The function takes the form:

```
MK.test(folder, sig.level = 5)
```

The help file can be found with the following code:

```
> ?MK.test
```

The function takes the following inputs:

Arguments	Description
<code>folder</code>	A path to a folder containing 2 sub folders. One must be called 'Inputs' and one must be called 'Outputs'. The 'Inputs' folder must contain .AM files for each gauge for which analysis is required. The 'Outputs' folder is where the 'Mann-Kendall' sub folder will be created, which is the location to which results files and plots will automatically be written.
<code>sig.level</code>	The significance level as a percentage for interpretation of results in the output .csv file (for example, input 5 to get an interpretation of whether trend is significant at the 5% significance level). 5 is the default setting for this parameter.

The function can be run as follows. The progress bar can be seen to update after each gauge has been analysed and the total processing time is also stated. If there are any warnings from running the function, they will also be stated in the console.

```
> folder <- "C:/Non-stationarity/Data"  
> MK.test(folder, sig.level = 5)
```

The outputs from running the function are as follows:

Output filename/format	Description
<i>Mann-Kendall results.csv</i>	A .csv file containing gauge information and results. Missing years are those in the middle of the AMAX series, not at the start or end, as there is no way of automatically determining that these are missing. The number of years reported is the number of years used in the analysis, excluding rejected years (that is, those marked as rejected in the .AM file) and rows with missing flow data (if they exist).
<i>gauge_number (gauge name – if exists) trend plots.png</i>	One plot per gauge containing 2 graphs. The first shows a time series of the AMAX data (black lines) and smoothed data (red line), using locally-weighted polynomial regression (LOWESS). The second shows the autocorrelation function. An example of the plot is shown in Figure 3-1, where in the second plot, the x-axis is the lag for each autocorrelation estimate and the y-axis is the estimate.

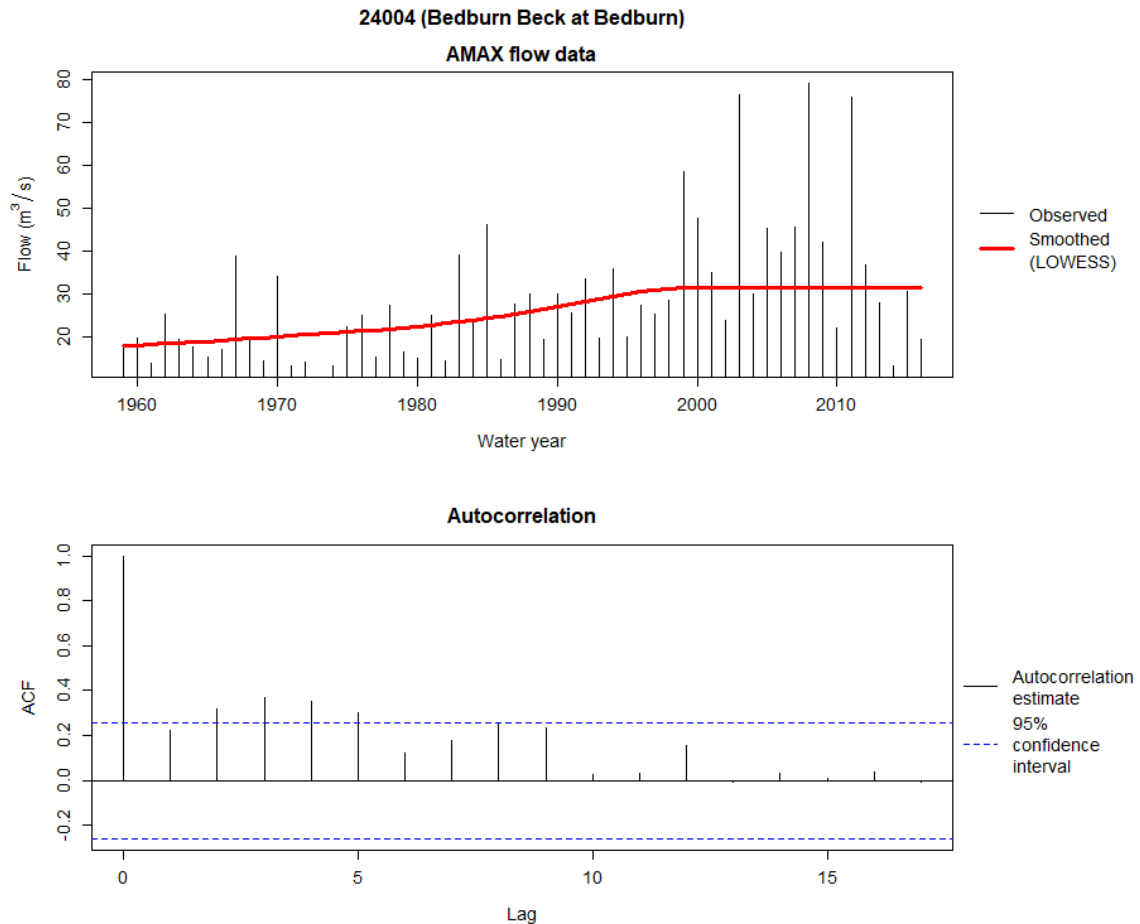
An explanation of the results from the Mann-Kendall test can be found in <https://cran.r-project.org/web/packages/trend/vignettes/trend.pdf>. The Mann-Kendall Z score (MKZ) is a standardised version of the results, so can be used to compare stations. Positive values of MKZ indicate increasing trends and negative ones indicate decreasing trends. The 2-sided p-value can be used to determine whether the results are statistically significant at a given significance level (as specified by the user) and the interpretation is reported in the next column. At a 5% significance level, this means that if the absolute value of MKZ is greater than 1.96 (or if the p-value is less than 0.05) then the null hypothesis (H0) of no trend is rejected. The conventional approach is then to (provisionally) accept a single alternative hypothesis H1, that is, that there is a trend. For other selected significance levels, the absolute values of MKZ must be greater than the following critical values to reject H0:

- 10% significance level: 1.645
- 2% significance level: 2.33
- 1% significance level: 2.58
- 0.1% significance level: 3.29

Autocorrelation plots can be used to check randomness in data. This is determined by computing autocorrelations for data values at a range of time lags. You can find further information at <https://www.itl.nist.gov/div898/handbook/eda/section3/autocopl.htm>.

This therefore enables investigation into whether there is significant lag-1 serial correlation in the AMAX time series (that is, whether the time series is dependent on its past), because the Mann-Kendall test requires serial independence of data. For example, if positive autocorrelation is present, it can increase the probability of detecting trends that do not exist. Therefore, other methods may be more suitable for assessing the significance of trend in the data, such as a modified version of the Mann-Kendall test (not available in the `nonstat` package).

Figure 3-1: Example of the graphical output of the `mk.test` function



The autocorrelation has been estimated at a range of lags (water years, given that AMAX data are being used). Be aware that missing years are not taken into account (that is, the function assumes that the values are for consecutive years), so the estimates may not be entirely correct, particularly if there are lots of missing years throughout the time series. The lag-1 autocorrelation indicates the correlation between values in the time series and their preceding value and lag-2 indicates the correlation between values and the value 2 water years before. The Bedburn Beck example above shows autocorrelations for lags up to 17 years.

The ACF at lag-0 is always 1. High values of autocorrelation indicate strong correlation between values at given lags. For time series where the current value is closely related to its preceding value, high values of autocorrelation would be expected for the shorter lags. For time series showing a periodic pattern, for example, where the current value is closely related to the observation 2 before it, there would a high value of autocorrelation for lag-2. For time series with no clear patterns, the values of autocorrelation should be relatively small. The blue, horizontal dashed lines on the plot represent lag-wise 95% confidence intervals centred around zero.

Autocorrelation estimates at a given lag would be statistically significant at a 5% significance level if they are outside these lines (compared to a null hypothesis of no autocorrelation).

3.2 Pettitt.test

This function carries out Pettitt's test, which is designed to detect a sudden change in the mean of a time series. It is a non-parametric test, that is, it makes no assumption about the distribution followed by the data. Results are automatically written out to the 'Outputs/Pettitt' folder. The 'Pettitt' folder will be automatically generated if it does not already exist.

The function takes the form:

```
Pettitt.test(folder, sig.level = 5)
```

The help file can be found with the following code:

```
> ?Pettitt.test
```

The function takes the following inputs:

Arguments	Description
folder	A path to a folder containing 2 sub folders. One must be called 'Inputs' and one must be called 'Outputs'. The 'Inputs' folder must contain .AM files for each gauge for which analysis is required. The 'Outputs' folder is where the 'Pettitt' sub folder will be created, which is the location to which results files and plots will automatically be written.
sig.level	The significance level as a percentage for interpretation of results in the output .csv file (for example, input 5 to get an interpretation of whether the change point is significant at the 5% significance level). 5 is the default setting for this parameter.

The function can be run as follows. The progress bar can be seen to update after each gauge has been analysed, along with a statement that an image is being saved, and the total processing time is also stated. Any warnings will also be stated in the console.

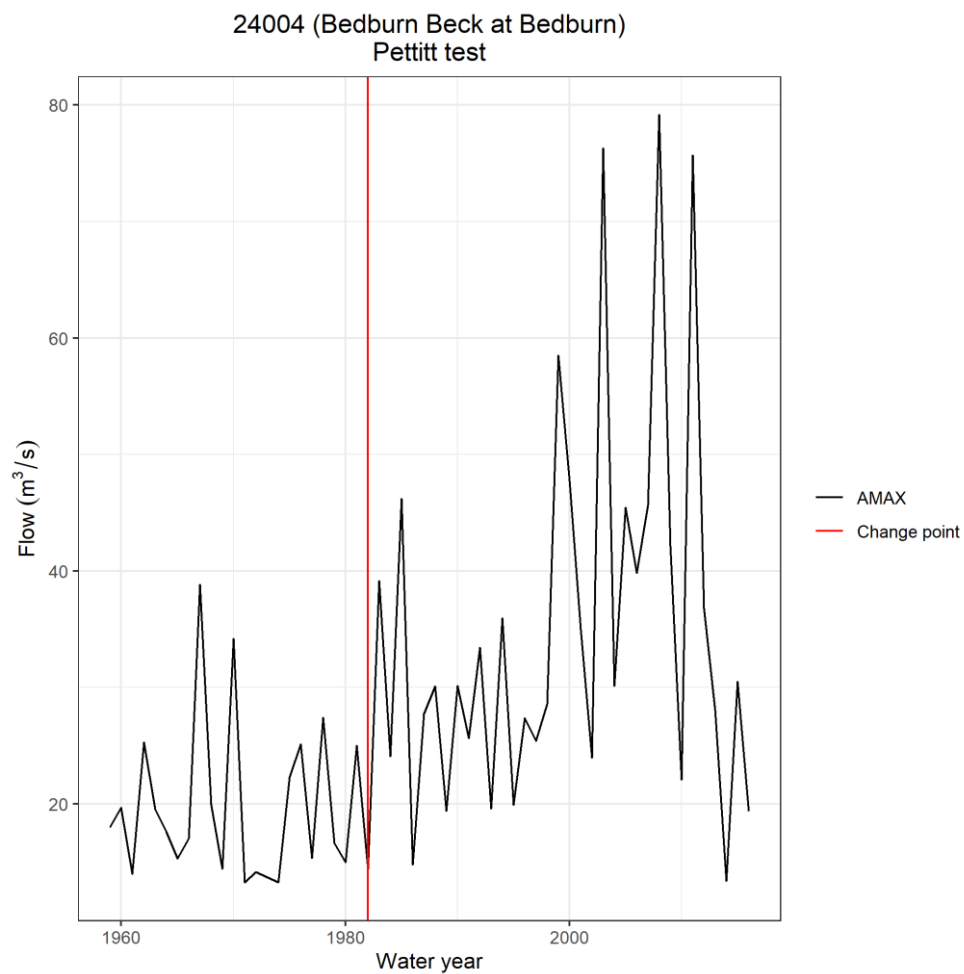
```
> folder <- "C:/Non-stationarity/Data"  
> Pettitt.test(folder, sig.level = 5)
```

The outputs from running the function are as follows:

Output filename/format	Description
<i>Pettitt results.csv</i>	A .csv file containing the results for all gauges. The Pettitt test always detects a change point and the year of the detected change point is given. The p-value can be used to identify whether the change in the mean AMAX flow from the period before to the period after the change point is significant. In this case, the null hypothesis, H_0 , is that there is no difference between the means of the earlier and later portions of each

	<p>AMAX flow series. If the p-value is less than the user-specified significance level (for example, 5% means that the p-value should be compared with 0.05), then H_0 is rejected. The conventional approach is then to (provisionally) accept a single alternative hypothesis H_1, that is, that there is a sudden change. The interpretation is provided in the file. The absolute and percentage change in the mean between the periods before and after the change point is provided, as well as the direction of the change (positive or negative).</p>
<p><i>gauge_number</i> (<i>gauge name – if exists</i>) <i>Pettitt results.png</i></p>	<p>A plot is also generated for each gauge showing the AMAX data with the detected change point shown as a vertical line, if it is significant at the specified significance level. An example plot is shown in Figure 3-2.</p>

Figure 3-2: Example of the graphical output of the `Pettitt.test` function



3.3 PELT.test

One limitation of Pettitt's test is that it can only detect a single change point. Additionally, it has been criticised for its tendency to classify gradual trends as sudden changes (Rougé and others, 2013). An alternative is the PELT (Pruned Exact Linear Time) test. As with most change point algorithms, PELT (Killick and others, 2012) tries to find the optimal segmentation in a time series. The method avoids identifying too many unrealistic change points. PELT is implemented through the `changepoint` R package (Killick and Eckley, 2014) and allows one or more change(s) in mean and variance of the AMAX flow data to be detected. This function is restricted to 2 change points; results from national analysis found no gauges where more than 2 change points were detected). The approach requires an assumption that the data follow a known distributional form. Since the flood time series are not generated from any known distribution, a log transformation is applied to transform the data to approximate normality (Box and Cox, 1964). In effect, it is assumed that the AMAX flows follow a log-normal distribution, which has been used in previous studies (for example, Prosdocimi and others, 2014). It is worth checking that the log-normal assumption is reasonable. You can do this by checking the normal Q-Q plot that is output from the package. Points close to the diagonal line indicate a reasonable log-normal fit. A minimum segment length is required, which prevents additional overfitting and false positive changes at short timescales. Therefore, 10 years has been chosen here as a suitable length for local stationarity pre- and post-change. This means that the record length of the input data must be at least 20 years (twice the minimum segment length). If the record length of your data is shorter than this, the function won't work and the gauge will be listed in a spreadsheet in the 'Outputs' folder: 'Gauges with record lengths too short to be analysed.csv'.

The function takes the form:

```
PELT.test(folder)
```

The help file can be found with the following code:

```
> ?PELT.test
```

The function takes the following inputs:

Arguments	Description
<code>folder</code>	A path to a folder containing 2 sub folders. One must be called 'Inputs' and one must be called 'Outputs'. The 'Inputs' folder must contain .AM files for each gauge for which analysis is required. The 'Outputs' folder is where the 'PELT' sub folder will be created, which is the location to which results files and plots will automatically be written.

The function can be run as follows. The progress bar can be seen to update after each gauge has been analysed, along with a statement that an image is being saved, and the total processing time is also stated. Any warnings will also be stated in the console.

```
> folder <- "C:/Non-stationarity/Data"
> PELT.test(folder)
```

The outputs from running the function are as follows:

Output filename/format	Description
<i>PELT results.csv</i>	A .csv file containing the results for all gauges. The year(s) of the detected change point(s) are given (some gauges will have no detected change points). For each detected change point, the absolute and percentage change in the mean and standard deviation between the periods before and after the change point is provided, as well as the direction of the change (positive or negative).
<i>gauge_number (gauge name – if exists) Normal Q-Q plot.png</i>	<p>A normal Q-Q plot of the log-transformed data. This indicates whether the log-normal assumption is reasonable; points close to the diagonal line should indicate a reasonable log-normal fit. If the log-transformed data do not remotely follow a normal distribution then the PELT results may not be reliable. Figure 3-3 shows 2 normal Q-Q plots.</p> <p>The points are close to the line in (a) and therefore the log-normal assumption seems to be reasonable. In (b), however, the points are further from the line and it is possible that the log-normal assumption is not valid and the PELT results may not be reliable.</p>
<i>gauge_number (gauge name – if exists) PELT results.png</i>	<p>A plot is also generated for each gauge showing the AMAX data with any detected change point shown as a vertical line. An example is given in</p> <p>Figure 3-4.</p>

Figure 3-3: Example of normal Q-Q plots from the `PELT.test` function

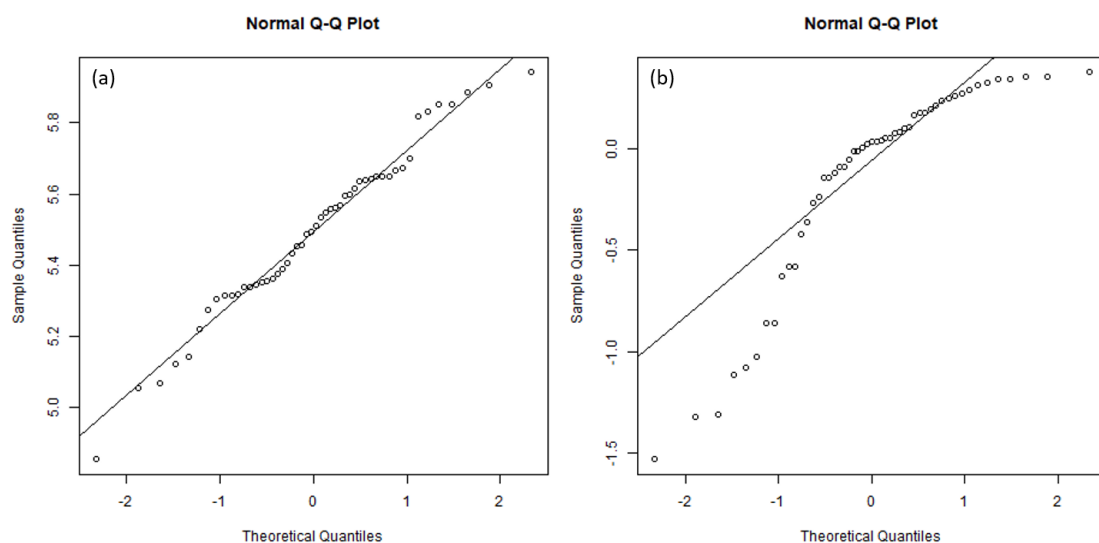
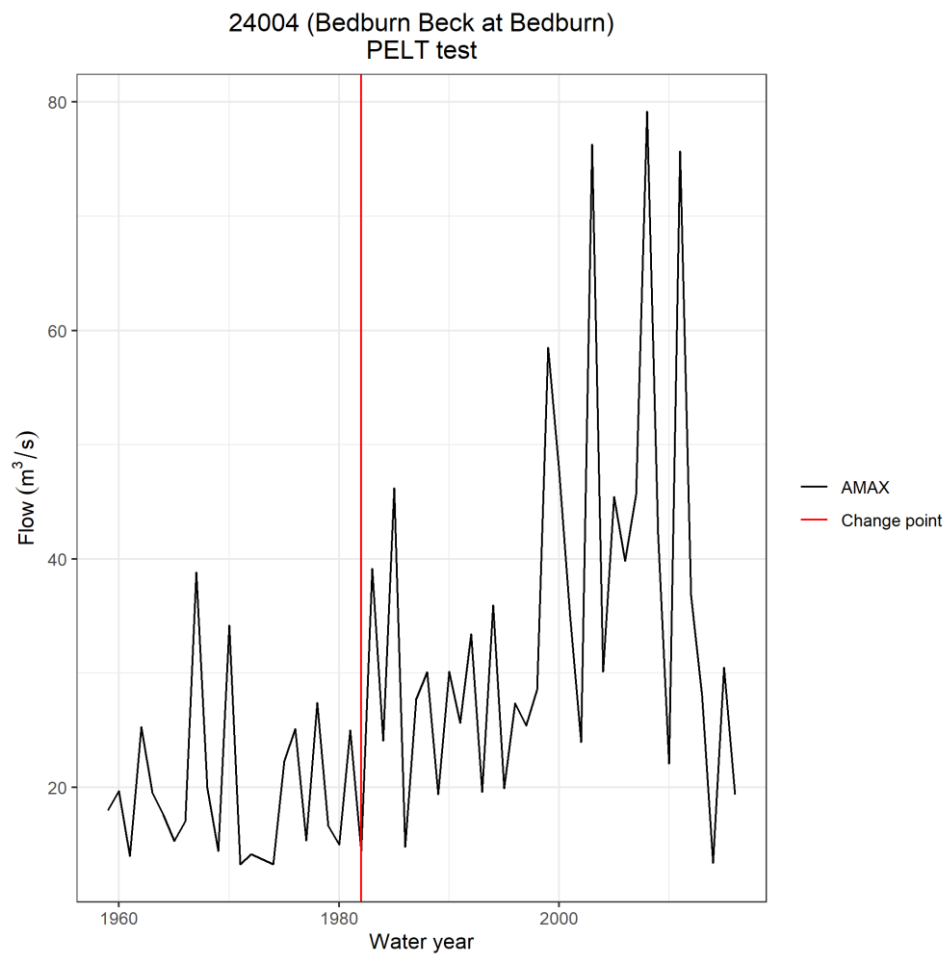


Figure 3-4: Example of a change point from the `PELT` test function



4 Fitting non-stationary models

Traditional statistical methods for flood frequency analysis assume that each AMAX flow in a gauge record represents an independent observation from the same probability distribution as all the other flow values (that is, observations of past flood events are assumed to represent the behaviour of future events). If this probability distribution is not constant over time (non-stationary), the peak flows are not identically distributed and so the assumption cannot be made. Therefore, if non-stationarity is evident in the AMAX flow data, then design estimates derived from traditional flood frequency analysis may be over or underestimated.

To model non-stationarity, covariates can be incorporated into the parameters of statistical models used in flood frequency analysis. Time (the year in which a flood happened) may seem to be an obvious choice for a covariate, relating changes in flood frequency to time. However, time is a proxy for some changing physical process that is causing change in flood frequency. An alternative option would be to use physically based covariates, such as rainfall or a climate index such as the North Atlantic Oscillation. Two reasons are sometimes given for this. The first is that physical covariates help remove some of the year-to-year variability in AMAX flows, enabling better identification of time-based trends and better fit of statistical distributions (Prosdocimi and others, 2014). The second is that they provide a more physically meaningful model of non-stationarity. Some physical covariates may open up the prospect of predicting the future evolution of the flood frequency curve (Sraj and others, 2016). They may not directly represent the physical processes that cause floods, but represent a potential step forward from the simplistic approach of modelling non-stationarity as a change over time.

This package has an option to use just time as a covariate in a statistical distribution, as well as an option for using physical covariates. The former is a simpler option, and can be interpreted as a trend test. Either the generalised extreme value (GEV) or generalised logistic (GLO) distributions can be used. The example code in this guidance uses the GEV distribution, but the GLO could be used by changing the `dist` parameter in relevant functions. All model fitting in this package is carried out using functions from the `texmex` package and all confidence intervals are calculated using bootstrapping. Given that bootstrapping involves an element of generating random numbers when sampling flows from the model (and when resampling covariates, where applicable), the seed has been set so that it effectively reuses the same random numbers each time the code is run. This means that results can be reproduced even though there is an element of randomness. The fitting function is set to use 2 cores when bootstrapping.

Typically, flood risk estimates are defined using return periods, but in non-stationary models, this becomes more complicated. This package produces a range of outputs that allow non-stationary models to be used, considering the different ways of defining risk.

Some key definitions for this section are:

Return level - the value (in this analysis, the river flow) associated with a particular return period or probability. Hydrologists often refer to this as the design flood, or as a point on the flood frequency curve. This is used in the practitioner guidance rather than return level.

Conditional return level – the return level that would be expected if the covariates had a particular combination of values (that is, conditional on those values). For instance, if the covariate was the water year, there could be a return level conditional on the water year being 2019 to 2020. This is referred to as a ‘conditional flow estimate’ in the practitioner guidance.

This is therefore suitable to use given a model with time covariates, as one could use the conditional return level given the current year.

It is less useful to think the same way when using physical covariates, as for example, the 100-year return level given the 2019 rainfall amount is the expected return level under the (clearly hypothetical) conditions that the rainfall always takes the value observed in 2019. In this case, an alternative approach would be to define a marginal return level, as used in Eastoe and Tawn (2009).

Marginal return level – the return level corresponding to the averaged encounter probability, where averaging is over covariates in a period of interest. Specifically, it is the return level where the average encounter probability is equal to the reciprocal of the return period. (An encounter probability is the probability of an event occurring in a given number of years.) The practitioner guidance refers to this as an ‘integrated flow estimate’.

Present-day marginal return level – the present-day expected return level for a particular exceedance probability, but not conditional on any particular value of a covariate such as annual rainfall. It is estimated by calculating the marginal return level, setting the time covariate to be the ‘current’ water year (last year in the flow and covariate data), and averaging over the full observed distribution of the physical covariates. This is a special case of the ‘single-year integrated flow estimate’ described in the practitioner guidance.

4.1 Fitting models where time is the only covariate

4.1.1 `fit.time`

The `fit.time` function fits stationary and non-stationary distributions (GEV or GLO) to AMAX flow data for all gauges provided by the user in an 'Inputs' folder. The 3 non-stationary models have a varying location parameter, a varying scale parameter and varying both location and scale parameters. The parameters vary with a covariate. This function uses time as a covariate and automatically calculates a value for the covariate to associate with each AMAX flow in the input data for the non-stationary distribution fitting, that is, the user does not need to supply the time covariate. The time covariate is an index of water year and the values are centred and scaled.

The results can be interpreted as a parametric trend test; there may be some catchments where this indicates non-stationarity, but where no trend is detected by non-parametric trend tests, such as the Mann-Kendall test.

Parameters of the fitted distributions and the results from frequency analysis are automatically written out to .csv files in the '*Outputs/Time covariate models*' folder. The '*Time covariate models*' sub folder will be automatically generated if it does not exist.

The function takes the form:

```
fit.time(dist, folder, RPs, CI = 90)
```

The help file can be found with the following code:

```
> ?fit.time
```

The function takes the following inputs:

Arguments	Description
<code>dist</code>	The statistical distribution to use in the fitting ("GEV" or "GLO").
<code>folder</code>	A path to a folder containing 2 sub folders. One must be called 'Inputs' and one must be called 'Outputs'. The 'Inputs' folder must contain .AM files for each gauge for which analysis is required. The 'Outputs' folder is where the 'Time covariate models' sub folder will be created, which is the location to which results files and plots will automatically be written.
<code>RPs</code>	The return periods for which frequency analysis is required.
<code>CI</code>	The required confidence interval for the results as a percentage, for example, 90 for 90% is the default.

The function can be run as follows. The progress bar can be seen to update after each gauge has been analysed and the total processing time is also stated. If there are any warnings from running the function, they will also be stated in the console.

In the example below, all the arguments of the function are set to sensibly named objects and then the outputs of the function are saved to an object called 'GEVfits'. It could be called anything you choose. It's used with some of the later functions.

```
# Set up data to use as inputs to the function

> dist <- "GEV"
> folder <- "C:/Non-stationarity/Data"
> RPs <- c(2, 10, 20, 50, 100)
> CI <- 90

# Run the fitting function, saving the outputs to an object
# called 'GEVfits'.

> GEVfits <- fit.time(dist, folder, RPs, CI)
```

The function fits a GEV or GLO distribution using maximum likelihood estimation (MLE) for a stationary case and 3 non-stationary cases (varying location, varying scale and varying both location and scale parameters). The scale parameter is log-transformed to ensure positivity for all possible covariate values.

The outputs from running the function are as follows:

Output filename/format	Description
<i>gauge_number</i> (<i>gauge name – if exists</i>) <i>GEV GLO results.csv</i>	A results spreadsheet containing the conditional return levels (flow estimates) from the fitted models. The first columns show the AMAX flow data used in the analysis. The 'WaterYearIndex' column is the automatically generated data representing time and the 'ScaledWaterYearIndex' column is the centred and scaled version of this, which is used as a covariate to represent time in the model fitting. For each model, the 'Fit' column is the type of fit, and the subsequent columns are the results from that fitted model (including the user-specified confidence intervals). The results from fitting a stationary distribution are the same for every year, but for the non-stationary cases, the results vary with time (for example, the 100-year flow in 1960 may be different from the 100-year flow in 2015), because they are conditional on a given year. Note that the results estimates here are taken from the median of the bootstrap sample and the confidence intervals are taken from the relevant quantiles of the bootstrap sample.
<i>GEV GLO distribution parameters and fits.csv</i>	A spreadsheet containing the parameters from the distribution fitting is also produced. There is one row per model, so 4 rows per gauging station. The first columns provide information about each station. Missing years are those in the middle of the AMAX series, not at the start or end, as there is no way of automatically determining that these are missing. The number of years reported is the number of years used in the analysis, excluding rejected years and rows with missing flow data (if they exist). The 'Fit' column gives the model type that corresponds to the parameters in that row. The next columns are the parameters of the fitted distribution.

	<p>For a stationary fit, the location, $\ln(\text{scale})$ and shape parameters are provided, along with their standard errors. The scale parameter (σ) can be obtained by taking the exponential of the value provided (which is $\ln(\sigma)$ or ϕ). For the non-stationary fits, the outputs change for the varying parameter(s). If the location parameter is varying, the location parameter (μ) is replaced with μ_0 and μ_1 (mu0 and mu1 in the column names). If the scale parameter is varying, the $\ln(\text{scale})$ parameter (ϕ) is replaced with ϕ_0 and ϕ_1 (phi0 and phi1 in the column names), as described in the following equations where the parameters vary with time, t:</p> $\mu(t) = \mu_0 + \mu_1 t,$ $\ln(\sigma(t)) = \phi(t) = \phi_0 + \phi_1 t$ <p>Therefore, if the parameter is varying, the column names should be read as mu0 and phi0, but if the parameter is not varying, they should be read as location and $\ln(\text{scale})$.</p> <p>The AIC and BIC are the Akaike Information Criterion and the Bayesian Information Criterion and represent the goodness of fit of the model. They are both methods of selecting the best fitting model for each station and establish a trade-off between the goodness of fit and the simplicity of the model. The lowest value indicates the best fit. The BIC tends to give preference to simpler models. It is worth noting that there can be small differences in the AIC or BIC value between different models for the same catchment and yet the models themselves can produce quite different results. It is therefore recommended that the results are assessed alongside diagnostic plots and other methods of judging goodness of fit to determine which model is most suitable.</p> <p>Another approach to find best fitting models is to use likelihood ratio testing. This is a preferred approach when comparing a small number of candidate models, when the likelihood ratio can be calculated for each nested pair of models. It is impractical when comparing hundreds of models, but it is suitable for this case where there are only 4 models to compare. The last 3 columns therefore identify the best fit per gauge based on the AIC value, the BIC value and likelihood ratio testing.</p>
<p><i>gauge_number</i> (<i>gauge name – if exists</i>) <i>GEV GLO</i> <i>diagnostic plots – fit type.png</i></p>	<p>This contains 3 plots (see example in Figure 4-1). The first is of the observed AMAX data, the second is a probability-probability (or P-P) plot and the third is a quantile-quantile (or Q-Q) plot. Residual P-P plots and residual Q-Q plots are specific instances of P-P and Q-Q plots. They are used because there are covariates in the models, so for a given data set (flow record) (with accompanying covariates), there is no such thing as 'the model quantiles' since each set of covariates defines a different fitted GLO or GEV distribution. Therefore, a P-P or Q-Q plot cannot be constructed for the original data, as they don't come from a common distribution. Instead, the data points are first transformed to a common distribution by constructing model residuals using the fitted model parameters for each data point. These residuals do come from a common distribution and therefore meaningful P-P and Q-Q plots can be constructed. A well-fitting model will have points lying close to the unit diagonal on both the P-P and Q-Q plots.</p> <p>The plot title contains information about the type of fit, that is, stationary or non-stationary. Information about covariates included in the model is in the form of the regression equation of covariates used in the location and scale parameters. For this <code>fit.time</code> function:</p> <ul style="list-style-type: none"> if the model is stationary, then no covariates are included, so the equation is just '~1' for both the location and scale parameters because they are constant

	<ul style="list-style-type: none"> • if the model is non-stationary with only the location parameter varying with time (that is, <code>ScaledWaterYearIndex</code>, the covariate that is automatically-generated within the package to represent time), then the equation for the location parameter is '$\sim 1 + \text{ScaledWaterYearIndex}$'. The scale parameter stays as '~ 1' as this is not varying with a covariate • if the model is non-stationary with only the scale parameter varying with time (that is, <code>ScaledWaterYearIndex</code>), then the equation for the scale parameter is '$\sim 1 + \text{ScaledWaterYearIndex}$'. The location parameter stays as '~ 1' as this is not varying with a covariate • if the model is non-stationary with both the location and scale parameters varying with time (that is, <code>ScaledWaterYearIndex</code>), then the equation for both is '$\sim 1 + \text{ScaledWaterYearIndex}$'.
<i>'gauge_number (gauge name – if exists) GEV GLO diagnostic plots – Q-Q plots for all models.png'</i>	The Q-Q plots for all 4 models are also provided for easy comparison (see example in Figure 4-2).
<i>GEV GLO diagnostic plot warnings and errors.csv</i>	If there were any warnings or errors produced when generating the diagnostic plots, then they will be stated in here. There is one column for warnings and one for errors for each type of model. It will be empty if there were no warnings or errors.
<i>GEV GLO fit time function warnings and errors.csv</i>	If there were any warnings or errors produced within the function used to fit the distributions and calculate return levels, they will be stated in here.
<i>GEV GLO model fitting warnings and errors.csv'</i>	If there were any warnings or errors produced during the actual model fitting, they will be stated in here. There is one column for warnings and one for errors for each type of model.

It's important to note that if there are any errors, then any results from the process which caused the error will be missing.

It is possible that the message 'Ratio of bias to standard error is high' will appear in the RStudio console window during the model fitting. This is an indication of the variability of the bootstrapped estimates and can be ignored in this instance.

Once this function has been run, there are 2 functions that can be used for plotting the results. These are `plot.selected.distributions` and `plot.all`.

Figure 4-1: Example diagnostic plot from the `fit.time` function: Observed AMAX data, probability-probability (or P-P) plot and a quantile-quantile (or Q-Q) plot.

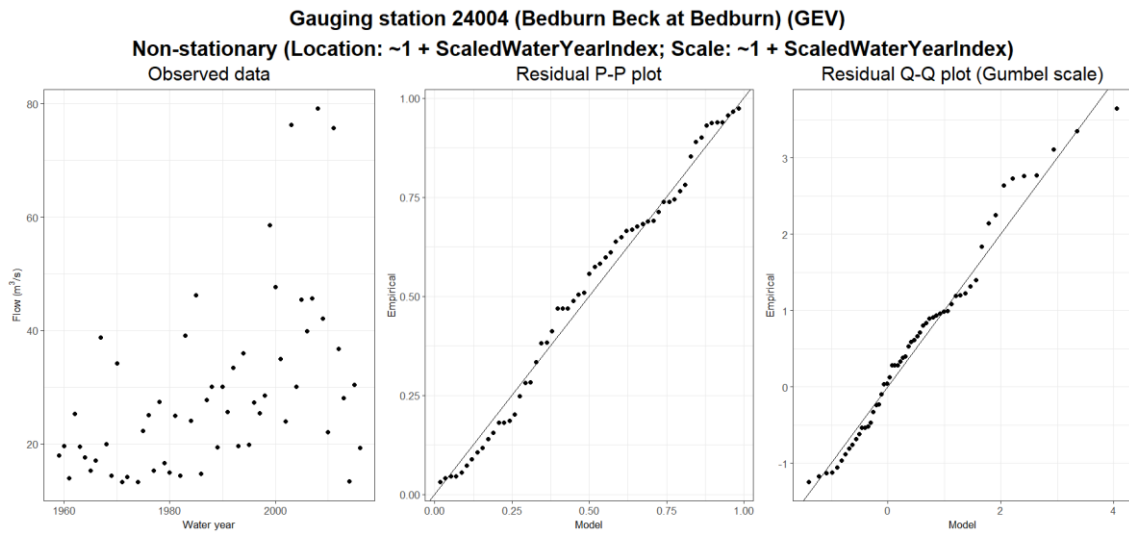
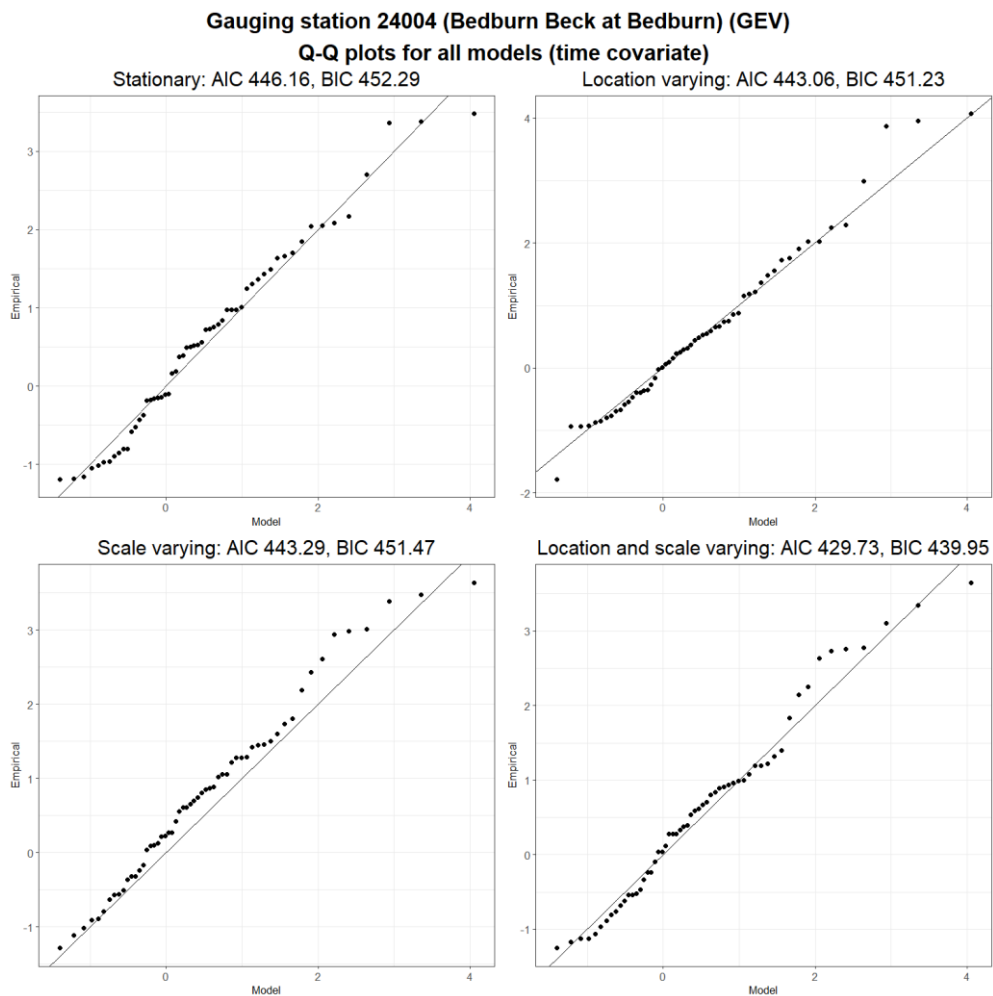


Figure 4-2: Example of Q-Q plots for all 4 models (time as covariate) from the `fit.time` function



4.1.2 plot.selected.distributions

This function plots the results from the fitted stationary distribution and one of the fitted non-stationary distributions selected by the user (for example, the best fitting non-stationary model based on likelihood ratio testing) for one user-selected gauge using the outputs from `fit.time` for a range of return periods. It also plots the confidence interval for one return period on a second plot. These plots are automatically written out to a .png file in the 'Outputs/Time covariate models' folder.

The function takes the form:

```
plot.selected.distributions(folder, gauge, type, fitResults,  
RPs, CI.RP, CI = 90, plotTitle = NULL)
```

The help file can be found with the following code:

```
> ?plot.selected.distributions
```

The function takes the following inputs:

Arguments	Description
folder	A path to a folder containing 2 sub folders. One must be called 'Inputs' and one must be called 'Outputs'. The 'Outputs' folder is where the 'Time covariate models' sub-folder will have been created, which is the location to which the plots will automatically be written.
gauge	The number of the gauge to plot; this must be one of the .AM files to whose data a model was fitted using the <code>fit.time</code> function.
type	The type of non-stationary distribution to plot. Options are 'Varying location', 'Varying scale' and 'Varying location and scale'.
fitResults	The output from the fitting function (<code>fit.time</code>). This could either be from a saved location or will already be in the R environment if <code>fit.time</code> has just been run and saved to a named object.
RPs	The return periods to plot (the maximum number of return periods allowed is 5).
CI.RP	The return period for which a confidence interval is to be plotted. Must be in RPs.
CI	The required confidence interval as a percentage to plot for the results from the distribution fitting, for example, 90 for 90% is the default.
plotTitle	This is an optional parameter. The code will automatically generate a plot title, but this can be used if you wish to change the title of the output plot. To split a title over 2 lines use a backslash followed immediately by 'n' in the character string where you want the line break to be.

The function can be run as follows:

```
# Set up data to use as inputs to the function

> folder <- "C:/Non-stationarity/Data"
> gauge <- "24004"
> type <- "Varying location and scale"
> fitResults <- GEVfits
> RPs <- c(2, 10, 20, 50, 100)
> CI.RP <- 100
> CI <- 90

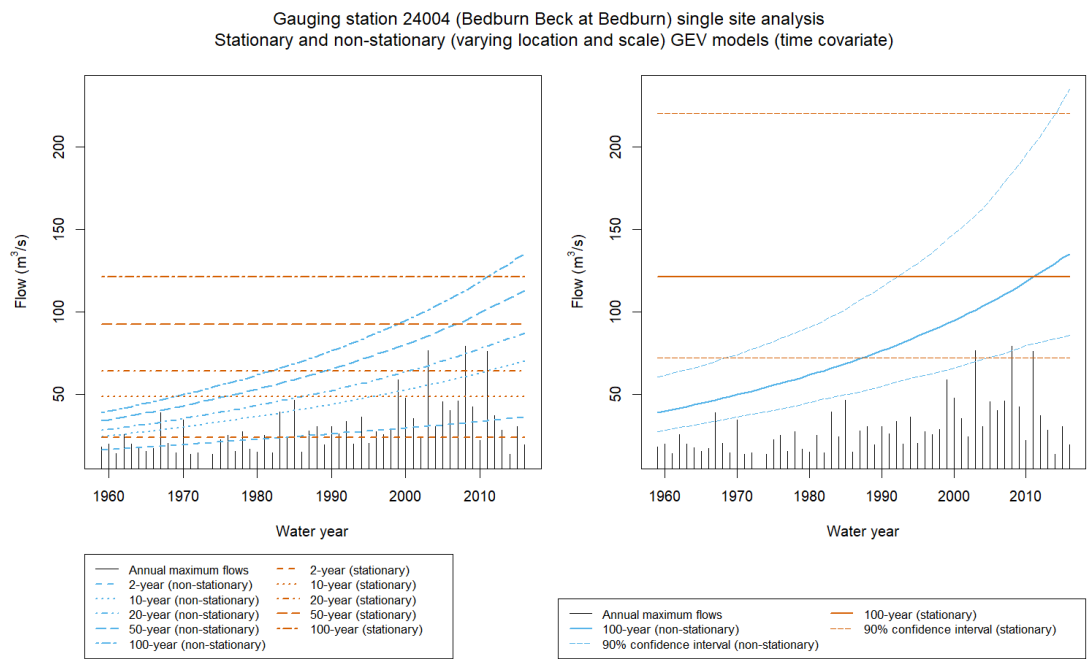
# Run the function

> plot.selected.distributions(folder, gauge, type, fitResults,
  RPs, CI.RP, CI)
```

The outputs from running the function are as follows:

Output filename/format	Description
<i>gauge_number</i> (<i>gauge name – if exists</i>) <i>GEV GLO</i> <i>user-specified distribution results</i> <i>and confidence interval plots.png</i>	An example plot is shown in Figure 4-3. The first plot shows the return levels from the fitted stationary model and the conditional return levels from one user-selected non-stationary model for specified return periods. In the example above, the varying location and scale parameters model results have been plotted, as stated in the plot title. The second plot shows confidence intervals for the return levels for one specified return period (in this example, for the 100-year return period). It can be seen that the non-stationary results are varying over time, with the 100-year flow increasing from 39 to 135 m ³ /s over the period of record.

Figure 4-3: Example of output plots from the `plot.selected.distributions` function



4.1.3 plot.all

This function plots the results from the fitted stationary and 3 non-stationary models (varying location, varying scale and varying location and scale parameters) with time as a covariate for a user-specified gauge. There will be one plot per specified return period. The plots will be automatically written out to a .png file in the 'Outputs/Time covariate models' folder.

The function takes the form:

```
plot.all(folder, gauge, fitResults, RPs, plotTitle = NULL)
```

The help file can be found with the following code:

```
> ?plot.all
```

The function takes the following inputs:

Arguments	Description
folder	A path to a folder containing 2 sub folders. One must be called 'Inputs' and one must be called 'Outputs'. The 'Outputs' folder is where the 'Time covariate models' sub folder will have been created, which is the location to which results files and plots will automatically be written.
gauge	The number of the gauge to plot; this must be one of the .AM files to whose data a model was fitted using the <code>fit.time</code> function.
fitResults	The output from the fitting function (<code>fit.time</code>). This could be either from a saved location or will already be in the R environment if <code>fit.time</code> has just been run and saved to a named object.
RPs	The return periods to plot.
plotTitle	This is an optional parameter. The code will automatically generate a plot title, but this can be used if you wish to change the title of the output plot. To split a title over 2 lines use a backslash followed immediately by 'n' in the character string where you want the line break to be. This should only be used when one return period has been specified, if you wish to manually include the return period value in the title. The automatically-generated title contains the appropriate return period.

The function can be run as follows:

```
# Set up data to use as inputs to the function

> folder <- "C:/Non-stationarity/Data"
> gauge <- "24004"
> fitResults <- GEVfits
> RPs <- c(2, 10, 20, 50, 100)

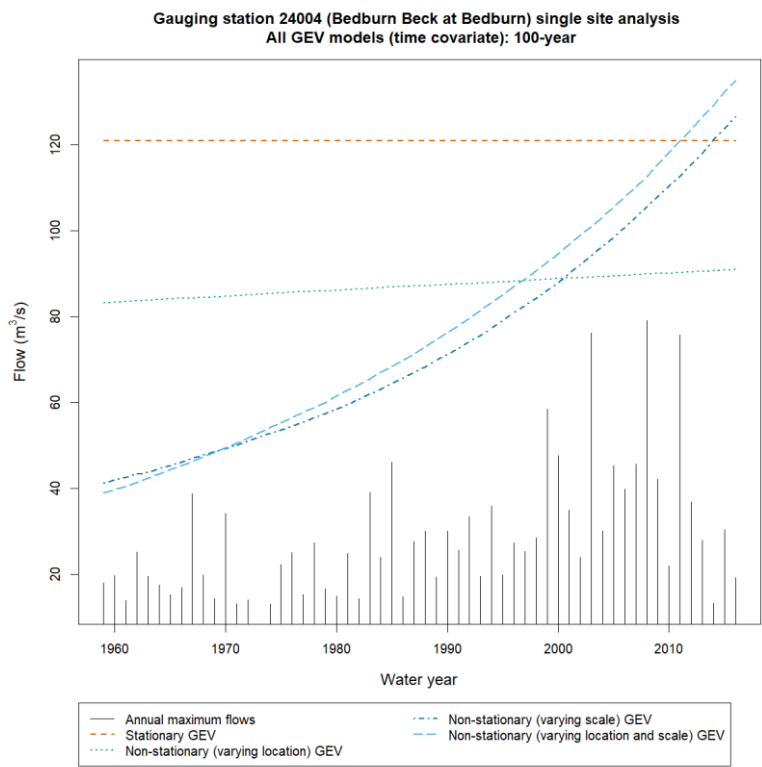
# Run the function

> plot.all(folder, gauge, fitResults, RPs)
```

The outputs from running the function are as follows:

Output filename/format	Description
<i>gauge_number</i> (<i>gauge name - if exists</i>) <i>GEV GLO all distribution plots</i> (<i>RPyear</i>). <i>png</i>	.png files (one per return period) containing plots showing results from all 4 fitted distributions. An example for the 100-year return period is shown in Figure 4-4.

Figure 4-4: Example of output plots from the `plot.all` function



For completeness and code efficiency, the following code gives an example of how you could use the `fit.time`, `plot.selected.distributions` and `plot.all` functions together:

```
# First set-up data to use as inputs to the fit.time function.

> folder <- "C:/Non-stationarity/Data"
> dist <- "GEV"
> RPs <- c(2, 10, 20, 50, 100)
> CI <- 90

# Now run the fitting function, saving the outputs to an object
# called 'GEVfits'

> GEVfits <- fit.time(dist, folder, RPs, CI)

# Once the fitting function has been run, you can run the
# plotting functions.

# Set up data to use as inputs to the
# plot.selected.distributions function, this will plot results
# from one of the fitted models. We'll use the RPs and CI we've
# already specified.

> gauge <- "24004"
> type <- "Varying location and scale"
> fitResults <- GEVfits
> CI.RP <- 100

# Run the function

> plot.selected.distributions(folder, gauge, type, fitResults,
  RPs, CI.RP, CI)

# Set up data to use as inputs to the plot.all function.
# Again we'll reuse some of the previous inputs, so no need to
# specify all the arguments (unless you want something
# different of course)

> plot.all(folder, gauge, fitResults, RPs)
```

In the example above the results from the `fit.time` function are saved in your R environment and can then be used in the plotting functions. If you wish to save the results so that they can be used later in the plotting functions (in a new R session), this can be achieved by saving the results as a `.RDS` file (a serialised version of the data which is saved with gzip compression by default).

```
# Save fitting results as a .RDS file

> saveRDS(GEVfits, "Outputs/Time covariate only/GEV fits from
  fit time function.RDS")

# To read this back into R at a later date:

> GEVfits <- readRDS("Outputs/Time covariate only/GEV fits from
  fit time function.RDS")
```

Note that the names of the input data do not have to match the names of the function arguments. In the above example, `fit.results` is being used as an input to the `fitResults` parameter. This could also be included as `fitResults = fit.results`, rather than just `fit.results`. You would need to specify the argument like this if the arguments were given in the wrong order in the function, for example,

```
plot.all(fitResults = fit.results, folder = folder, RPs = RPs,
gauge = gauge) would also work.
```

If you wish to plot results for more than one gauge at a time, it is possible to write some code to loop through each gauge. This would only be possible for the `plot.selected.distributions` function if the same type of non-stationary model is to be fitted for each gauge. The following code could be used to loop through the gauges.

```
# Example of looping through gauges for plotting
# Set parameters

> dist <- "GEV"
> folder <- "C:/Non-stationarity/Data"
> RPs <- c(2, 10, 20, 50, 100)
> CI <- 90
> CI.RP <- 100
> type <- "Varying location"

# Fit models

> GEVfits <- fit.time(dist, folder, RPs, CI)
> fitResultsGEV <- GEVfits

# List gauges to loop through

> gauges <- list.files(paste0(folder, "/Inputs"), pattern =
  ".AM", full.names = TRUE)

# Read in the metadata file in your 'Inputs' folder

> metadata <- read.csv(paste0(folder,
  "/Inputs/Gauge_metadata.csv"), stringsAsFactors = FALSE)

# Loop through gauges

> for (i in 1:length(gauges)){
  print(gauges[i])
  preprocessedData <-
    import.preprocess.data(gauges[i],
      metadata = metadata)

  if (!is.null(preprocessedData$res)){
    gauge <- preprocessedData$res$gaugeNumber

    plot.selected.distributions(folder, gauge, type,
      fitResultsGEV, RPs, CI.RP, CI)

    plot.all(folder, gauge, fitResultsGEV, RPs)
  }
}
```

4.2 Fitting models with physical covariates

The functions described in this section allow users to incorporate physical covariates into flood frequency analysis.

One important consideration is the effect of collinearity and confounding variables on results when including multiple covariates. Ideally, covariates included in a non-stationary statistical model should be independent, otherwise there may be computational issues and the fitted model may be difficult to interpret. When time is one of the covariates, it is recommended that this is achieved by detrending the physical covariates based on the time covariate. This can be done in the `fit.phys.cov` function by using the `detrend.cov` argument.

There are 2 steps to carry out analysis using physical covariates in this package. The first is fitting models (varying the parameters by a range of covariate combinations) to flow data for all gauges in the 'Inputs' folder. The second is obtaining return level estimates (flows) for each gauge individually using one of the fitted models for that gauge.

Unlike the time covariate, physical covariates need to be supplied by the user in a .csv file (see the example file `24004_covariates.csv`). This should be included in the 'Inputs' folder with the flow data. The gauge number in the covariates file name must match the one in the .AM file, so if it has been amended in the .AM file, for example, `24004x`, then the covariates file should be called `24004x_covariates.csv`. A value of the covariate must be associated with every AMAX flow (for example, a value of annual or seasonal rainfall would be required for every water year in the flow record). The flow and covariate data should therefore cover the same time period, although the package will automatically trim both data sets to use only the overlapping time period. The file must contain a maximum of 7 physical covariates, otherwise the `fit.phys.cov` function will produce an error. Gauges without a covariate data file will be skipped.

The example file contains a range of covariates. The NAO and EA indices were obtained from NOAA (<https://www.cpc.ncep.noaa.gov/data/teledoc/telecontents.shtml>). The indices are calculated from anomalies that are standardised by monthly means and standard deviations calculated over a 1950 to 2000 baseline period.

In all cases, annual values of the covariates were calculated using water years. Each rainfall covariate was centred and scaled by subtracting the mean of the covariate time series from each observation in the record and then dividing the resulting values by the standard deviation of the time series.

This can be carried out using the 'scale' function in R. The other covariates were already standardised in a similar way. It is recommended that all covariates are centred and scaled.

4.2.1 fit.phys.cov

As with the `fit.time` function, the `fit.phys.cov` function fits stationary and non-stationary GEV or GLO distributions to AMAX flow data for all gauges provided by the user in an 'Inputs' folder. The nonstationary models have varying location and/or scale parameters that can vary with one or more physical covariate(s) and/or time. Outputs from the function are automatically written out to the 'Outputs/Physical covariates models' folder. The 'Physical covariates models' sub folder will be automatically generated if it does not exist. Time (a centred and scaled index of water year calculated within the function and given the name "ScaledWaterYearIndex") is automatically included as a covariate alongside the physical covariates, although the function fits a range of models with various combinations of covariates, so time will not be included in every model fitted.

The function takes the form:

```
fit.phys.cov(dist, folder, detrend.cov = "ScaledWaterYearIndex",  
option = 2, best_fit_method = "BIC")
```

The help file can be found with the following code:

```
> ?fit.phys.cov
```

The function takes the following inputs:

Arguments	Description
dist	The statistical distribution to use in the fitting ("GEV" or "GLO").
folder	A path to a folder containing 2 sub folders. One must be called 'Inputs' and one must be called 'Outputs'. The 'Inputs' folder must contain .AM files and .csv files containing the physical covariates data for each gauge for which analysis is required. The 'Outputs' folder is where the 'Physical covariates models' sub folder will be created, which is the location to which results files will automatically be written.
detrend.cov	If included (that is, not NULL), the function detrends all the covariates based on the covariate in this string, usually time (represented by "ScaledWaterYearIndex" in this function). This is recommended for models where a physical covariate and time may both be included in the location or scale parameter at the same time (options 1 and 2 below). The default is "ScaledWaterYearIndex" (that is, by default, all the physical covariates are detrended based on the time covariate), but this can be changed to NULL if you don't want to detrend the covariate data. It could also be set to any covariate included in the covariates input .csv file (the string here must match the name in the .csv file).
option	Must be 1, 2 (default) or 3. Option 1 fits models using all possible combinations of covariates supplied in the .csv file (this can produce tens of thousands of fitted models if many covariates are considered, for example, ~65,000 models for 7 physical covariates plus time); Option 2 fits models using combinations which only include one physical covariate and/or time (this would produce 88 fitted models

	for 7 physical covariates plus time); Option 3 fits models using combinations which include one covariate at a time in the location and/or scale parameter.
best_fit_method	Optional. Method of automatically selecting a best fit model that can be used in the next step if desired. Must be "AIC" or "BIC" (default).

The function can be run as follows:

```
# Set up data to use as inputs to the function

> dist <- "GEV"
> folder <- "C:/Non-stationarity/Data"

# Run function and save outputs to an object named
# 'physCovModelFits'

> physCovModelFits <- fit.phys.cov(dist, folder, detrend.cov =
  "ScaledWaterYearIndex", option = 2, best_fit_method = "BIC")
```

The function fits a GEV or GLO distribution using maximum likelihood estimation for a stationary case and a number of non-stationary cases (depending on the option selected). The scale parameter is log-transformed to ensure positivity for all possible covariate values.

The outputs from running the function are as follows:

Output filename/format	Description
<i>gauge_number</i> (<i>gauge name</i> - <i>if exists</i>) GEV GLO fitting information (<i>physical</i> <i>covariates</i>).c sv	<p>A spreadsheet containing the information about each fitted model. There is one row per model. The locFit column gives the regression equation of covariates included in the location parameter for that model. The scale-Fit column gives the regression equation of covariates included in the scale parameter for that model. For example:</p> <ul style="list-style-type: none"> if the location or scale parameter is constant (that is, no covariate is included in it), then the equation for that parameter is '~ 1'. If the model is stationary, then no covariates are included in either parameter, so the equation is just '~ 1' for both the location and scale parameter because they are constant if the model is non-stationary with the location and/or scale parameter varying with time (that is, Scaled-WaterYearIndex, the covariate that is automatically-generated within the package to represent time), then the equation for the varying parameter/s is '~ 1 + ScaledWaterYearIndex' if the model is non-stationary with the location and/or scale parameter varying with a physical covariate (for example, annual rainfall, which may be called Rain_WY in the input covariates file and therefore also in the outputs), then the equation for the varying parameter/s is '~ 1 + Rain_WY' if the model is non-stationary with the location and/or scale parameter varying with time and a physical covariate, then the

	<p>equation for the varying parameter/s may be '$\sim 1 + \text{Rain_WY} + \text{Scaled-WaterYearIndex}$'</p> <p>Additional covariates are included in the equation as extra terms. If the location and/or scale parameter was varying with the maximum of 7 physical covariates allowed in the package, as well as time, then the equation would have 9 terms: '$\sim 1 + \text{Covariate1} + \text{Covariate2} + \text{Covariate3} + \text{Covariate4} + \text{Covariate5} + \text{Covariate6} + \text{Covariate7} + \text{ScaledWaterYearIndex}$'.</p> <p>The warnings and errors columns give any warnings or errors that occurred during the fitting of the model (this is more likely when option 1 is selected as there may be many covariates in one model). Models with warnings or errors should not be selected for calculating return levels. Therefore, it is important to check this spreadsheet before running the 'res.phys.cov' function. The AIC and BIC are as described for the models with only time as a covariate. This file also contains columns that rank the models in terms of AIC and BIC. If <code>detrend.cov</code> has been set, this is reported in the 'DetrendingCovariate' column, and the 'CovariatesOption' column records which option was selected when the function was run.</p> <p>The 'params' column gives the parameters of the fitted distribution. For a stationary fit, the location, $\ln(\text{scale})$ and shape parameters are provided. However, as for the models where time is the only covariate, for the non-stationary fits, the outputs change depending on the varying parameter(s). If the location parameter is varying, then the location parameter (μ) is replaced with a μ value for the intercept and each covariate. This is equivalent to using μ_0 to μ_n (as used in the <code>fit.time</code> function results where there was only one covariate and therefore there was only a μ_0 and a μ_1). If the scale parameter is varying then the $\ln(\text{scale})$ parameter (ϕ) is replaced with a ϕ value for the intercept and each covariate, equivalent to ϕ_0 to ϕ_n. This is described in the following equations where the parameters vary with a vector of covariates, x</p> $\mu(x) = \mu_0 + \mu_1 x_1 + \dots + \mu_n x_n$ $\ln(\sigma(x)) = \phi(x) = \phi_0 + \phi_1 x_1 + \dots + \phi_n x_n$ <p>where n is the number of covariates included in the parameter.</p> <p>For example, in the above equation, if covariate x_1 represents time and covariate x_2 is a measure of rainfall, μ_1 would express how floods change over time, once variations in rainfall are taken into account by μ_2 (Prosdocimi and others, 2014).</p>
<p><i>gauge_number</i> (gauge name – if exists) <i>GEV GLO</i> <i>model fits</i> (physical covariates).<i>RDS</i></p>	<p>This is for reference purposes, if any information from the fitting process needs to be accessed again in the future. It is a list containing information for each gauge analysed. For each gauge, the following information is saved: <code>dist</code> (distribution used - taken from the input argument), <code>model_fits</code> (described below), <code>all_data_raw</code> (the data used in the model fitting - not detrended), <code>gaugeName</code> (name of the gauge) and <code>detrend.cov</code> (the string from the input argument). Most of these are also returned from the function so that they can be used in the next step (see code example).</p> <p>The <code>model_fits</code> element is another list containing the following information: <code>best_model</code> (details of the best fitting model selected based on AIC or BIC, as specified by the user), <code>best_model_formula</code> (details of the formula used to fit <code>best_model</code>), <code>stat_model</code> (details of the fitted stationary model) and <code>all_models_metadata</code> (a metadata file associated with fitting models to all combinations of covariates selected by the user). The latter is the same as the .csv file described above.</p>

Once this function has been run, the `res.phys.cov` function can be run for each gauge to obtain return levels (flows) based on one of the fitted models, selected by the user.

4.2.2 `res.phys.cov`

The `res.phys.cov` function uses one of the fitted non-stationary models from the `fit.phys.cov` function to estimate conditional, marginal and present-day marginal return levels (as described in the ‘Fitting nonstationary models’ section of this document). Stationary return levels are also estimated based on the fitted stationary model. The function also outputs plots of the results. The following inputs are required:

The function takes the form:

```
res.phys.cov(folder, physCovModelFits, gauge, RPs, CI = 90,
model_type = "Best", locFit = NULL, scaleFit = NULL,
present_day)
```

The help file can be found with the following code:

```
> ?res.phys.cov
```

The function takes the following inputs:

Arguments	Description
<code>folder</code>	As for the <code>fit.phys.cov</code> function. The results files will automatically be written out to ‘Outputs/Physical covariates models’.
<code>physCovModelFits</code>	Output from <code>fit.phys.cov</code> function.
<code>gauge</code>	Number of gauge to analyse. Must have been used in the <code>fit.phys.cov</code> function and must match the number in the fitting information .csv file output from the <code>fit.phys.cov</code> function (so if the number has been edited then it should be the edited number, for example, 24004x rather than 24004).
<code>RPs</code>	The return periods for which frequency analysis is required. A maximum of 8 return periods can be plotted on the ‘model results’ plot - if more are listed here then only the first 8 will be selected for the plot. All return periods will be used for the frequency analysis and for the ‘encounter probability’ plot.
<code>CI</code>	The required confidence interval for the results as a percentage, for example, ‘90’ for 90% is the default.
<code>model_type</code>	Must be “Best” (default) or “User”. This is for the user to select whether to use the best fit model based on AIC or BIC (whichever they specified in the <code>fit.phys.cov</code> function) or a user-selected model to calculate return levels (and then plot). The selected model should ideally be a non-stationary model as results will automatically be written out for the stationary model. Some results will not be available for a stationary model.

locFit	This is only required if <code>model_type = "User"</code> . It is the regression equation of covariates to include in the location parameter (copy from the 'GEV GLO fitting information (physical covariates)' .csv file produced by the <code>fit.phys.cov</code> function for the gauge), e.g. <code>locFit = "~ 1 + NAO_SON + ScaledWaterYearIndex"</code> .
scaleFit	This is only required if <code>model_type = "User"</code> . It is the regression equation of covariates to include in the scale parameter (copy from the 'GEV GLO fitting information (physical covariates)' .csv file produced by the <code>fit.phys.cov</code> function for the gauge), for example, <code>scaleFit = "~ 1 + ScaledWaterYearIndex"</code> .
present_day	Indicates whether to calculate the present-day marginal return level. Set to <code>TRUE</code> to calculate the present-day marginal return level as well as the marginal return level, otherwise set to <code>FALSE</code> . This can only be <code>TRUE</code> if time ("ScaledWaterYearIndex") is included as a covariate in the selected model. If it is set to <code>TRUE</code> and "ScaledWaterYearIndex" is missing from the model covariates then it will automatically be reset to <code>FALSE</code> .

The function can be run as follows:

```
# Obtain return levels based on one of the fitted models
# Set inputs

> folder <- "C:/Non-stationarity/Data"
> gauge <- "24004"
> RPs <- c(2, 10, 20, 50, 100)
> CI <- 90

# Run function
> res.phys.cov(folder, physCovModelFits, gauge, RPs, CI,
  model_type = "Best", locFit = NULL, scaleFit = NULL,
  present_day = TRUE)

# Note that the physCovModelFits argument comes from the output
# of the fit.phys.cov function (see previous section)
```

The outputs from running the function are as follows:

Output filename/format	Description
<i>gauge_number</i> (<i>gauge name – if exists</i>) <i>GEV GLO model.RDS</i>	A .RDS file is returned which contains information relating to the analysis for the gauge (from the model used to estimate return levels to the results of the frequency analysis). This can be opened in R and referred to at a later date if required. Much of the information is also written out in a range of .csv files, as described below.
<i>gauge_number</i> (<i>gauge name – if exists</i>) <i>GEV GLO</i>	This contains the regression equations used in the location and scale parameters for the selected model (that is, which covariates have been included). See previous section for explanation of these equations.

<i>model formula.csv</i>	
<i>gauge_number (gauge name – if exists) GEV GLO conditional return levels.csv</i>	This is similar to the 'gauge_number (gauge name – if exists) GEV GLO results.csv' output from the fit.time function. The Flow column is the observed AMAX flow data and the WaterYear column is the water year associated with each flow. Other columns at the beginning show the covariate data used in the model fitting (detrended, if detrend.cov has been used). If time is included as a covariate in the selected model, then the WaterYearIndex column is the automatically-generated index representing time and the ScaledWaterYearIndex column is the centred and scaled version of this, which was actually used as the time covariate in the model fitting. The conditional return levels (and their confidence intervals) are given in the remaining columns (these are taken from the median of the bootstrap sample). These vary based on whichever covariates have been used in the model selected in this function for calculating return levels. As previously mentioned, given that multiple covariates may have been used, these results can be less useful than those derived from stationary models or the non-stationary models with only time as a covariate. For example, the return level may be conditional on a specific water year and rainfall amount.
<i>gauge_number (gauge name – if exists) GEV GLO marginal return levels.csv</i>	This provides the marginal return levels (and confidence intervals) for each return period. Unlike the conditional return levels, there is only one value per return period. It indicates the return level corresponding to a particular averaged exceedance probability (the encounter probability) during the entire period of record. For example, the 100-year marginal return level is the value expected to be exceeded once every 100 years assuming the covariate distribution in the observed period is unchanged. There may be a higher chance of that exceedance occurring in the second half of the record, if the conditional return levels are higher in that half. These are calculated from the fitted model, rather than the median of the bootstrap sample; the confidence intervals are from the relevant quantiles of the bootstrap sample. The same applies for the present-day marginal return levels below.
<i>gauge_number (gauge name – if exists) GEV GLO present-day marginal return levels.csv</i>	This provides the present-day marginal return levels (and confidence intervals) for each return period. This is similar to the marginal return levels, except that the time covariate has been set to be the last year in the input flow and covariate data (representing the present-day) and then the averaging is done over the full observed distribution of the physical covariates.
<i>gauge_number (gauge name – if exists) GEV GLO stationary return levels.csv</i>	This provides the return levels for each return period based on the stationary model. This is the usual definition of a return level (the flow associated with a particular return period).
<i>gauge_number (gauge name – if exists) GEV GLO model diagnostic plots.png</i>	This contains the diagnostic plots relating to the selected model. These should be checked to make sure that the selected model has a good fit, rather than just relying on the AIC or BIC value. The plots are the same as those from the fit.time function, but the titles will also contain physical covariates in the location and scale parameter equations, as appropriate.
<i>gauge_number (gauge name – if exists) GEV GLO</i>	This is a plot of the flow data plotted against all covariates in the input file (detrended, if detrend.cov has been used) and can be used to see whether flow is correlated with the covariates used. An example plot is given in Figure 4-5 .

<i>correlation plot.png</i>	
<i>gauge_number (gauge name – if exists) GEV GLO model results plot.png</i>	This is a plot of flow plotted against conditional return levels from the non-stationary model and stationary return levels from the stationary model for the observed period. An example can be seen in Figure 4-6 . The dots are the conditional return levels and these will move up and down each year, depending on the covariates. The stationary return levels are the dashed horizontal lines. If the selected model is stationary then only the stationary return levels will be plotted as conditional return levels will not be available.
<i>gauge_number (gauge name – if exists) GEV GLO encounter probability plot.png</i>	This is a plot of flow against encounter probability, in this case, the probability of the flow being exceeded over the full period of record. An example can be seen below. The plot shows both the stationary return levels ('Stationary flow estimates') and marginal return levels (referred to as 'Integrated flow estimates', as in the practitioner guidance), including a confidence interval for the latter. The observed AMAX flows are plotted as short black lines on the y-axis. If the selected model is stationary, then marginal return levels will not be available and therefore only the stationary results will be plotted, along with their confidence interval. The smoothness of the curve and range of results shown on the plot will depend on the number of return periods selected by the user because the plot uses the same stationary and marginal return levels that are output in the .csv files.

Figure 4-5: Example of output plot of flow data against all covariates in the input file from the `res.phys.cov` function




Figure 4-6: Example of output plot of flow against conditional return levels from the non-stationary model and stationary return levels from the stationary model for the observed period from the `res.phys.cov` function





Figure 4-7: Example of output plot of flow against encounter probability from the `res.phys.cov` function



Any warnings and errors associated with the analysis will be written out in the R console. Due to the use of bootstrapping to calculate confidence intervals, 200 (the number of bootstrap samples used) models are fitted using maximum likelihood estimation. There are therefore likely to be some errors and warnings associated with some of these models.

You may get messages stating that there were a number of warnings or errors when fitting the model for a particular bootstrap sample. This is for information only. You may then also get a message stating that there are errors or warnings in more than 10 of the bootstrap models for at least one return period. This means that there may be issues with at least 5% of the 200 bootstrap models used to find marginal or present-day marginal (present-day will be stated in the message if applicable) return level confidence intervals.

In addition, you may get a message stating the number of errors or warnings that occurred when calculating marginal return levels from the 200 models in order to

calculate the confidence intervals over all the required return periods. Bootstrap samples with errors or warnings will probably result in NAs and therefore not be used to calculate the confidence interval. There may also be additional errors from the model fitting that appear in the console, such as 'Error in solve.default(o\$hessian)' or 'Error in diag(o\$cov)'. These imply poor-fitting of the model. Therefore, if there are a large number of errors or warnings, it may be that less confidence should be placed in the values of the confidence interval. It is worth noting that non-convergence indicates that there is an issue somewhere in the model assumptions on which the estimation process is based, so it is useful to understand the context of the data and the model being fitted to it.

Information about the bootstrapping results, warnings and errors is also saved in the .RDS file under 'margRLsAllInfo'. If you do need to refer to this information, the explanation for each output can be found in the 'BootstrapMargRLs' function help file (this function is only used within other functions and therefore is not included in this document). The information includes the marginal return levels for each bootstrap sample and return period, so you can see where the problems are. It is possible that there may be some poorly-fitting bootstrap models that have been used to produce results for lower return periods but have resulted in NAs for higher return periods, in which case, it is likely that any results produced are not sensible. If this is a problem with lots of the bootstrap samples, then the confidence interval for each return period may be less accurate. Additionally, any models containing warnings or errors in the 'fitting information' .csv file that is output from the `fit.phys.cov` function should not be used. If any confidence intervals are coming out the same as the estimate, this also implies that there were problems with the model fitting and that the results should not be used.

For completeness and code efficiency, the following code gives an example of how you could use the `fit.phys.cov` and `res.phys.cov` functions together:

```
# Fit models using physical covariates

> folder <- "C:/Non-stationarity/Data"
> dist <- "GEV"

# Run function and save outputs to an object named
# 'physCovModelFits'

> physCovModelFits <- fit.phys.cov(dist, folder, detrend.cov =
  "ScaledWaterYearIndex", option = 2, best_fit_method = "BIC")

# Obtain return levels based on one of the fitted models
# Set inputs

> gauge <- "24004"
> RPs <- c(2, 10, 20, 50, 100)
> CI <- 90

# Run function

> res.phys.cov(folder, physCovModelFits, gauge, RPs, CI,
  model_type = "Best", locFit = NULL, scaleFit = NULL,
  present_day = TRUE)
```

5 Complete example script to run all functions

```
# Setup -----
# Install package dependencies

> install.packages(c('changeoint', 'trend'))

# Install the nonstat package from saved location

> path_to_file <- "C:/Non-stationarity/nonstat_1.0.0.tar.gz"
  install.packages(path_to_file, repos = NULL, type="source")

# Load the nonstat package

> library(nonstat)

# Set the path to the folder containing the 'Inputs' and
# 'Outputs' folders

> folder <- "C:/Non-stationarity/Data"

# Trend testing -----

> MK.test(folder, sig.level = 5)
> Pettitt.test(folder, sig.level = 5)
> PELT.test(folder)

# Flood frequency analysis (time only as a covariate) -----
# Run the fitting function

> dist <- "GEV"
> RPs <- c(2, 10, 20, 50, 100)
> CI <- 90
> GEVfits <- fit.time(dist, folder, RPs, CI)

# Run the plotting functions
> gauge <- "24004"
> type <- "Varying location and scale"
> fit.results <- GEVfits
> RPs <- c(20, 50, 100) # return periods to plot
> CI.RP <- 100 # return period for confidence interval plot
> plot.selected.distributions(folder, gauge, type, fit.results,
  RPs, CI.RP, CI)
> RPs <- c(2, 10, 20, 50, 100) # return periods to plot
> plot.all(folder, gauge, fit.results, RPs)

# Continued on next page...
```

```

# Continued from previous page...

# Flood frequency analysis (physical covariates) -----

# Run the fitting function

> detrend.cov <- "ScaledWaterYearIndex"
> option <- 2
> best_fit_method <- "BIC"
> physCovModelFits <- fit.phys.cov(dist, folder, detrend.cov,
  option, best_fit_method)

# Obtain return levels based on one of the fitted models

> gauge <- "24004"
> RPs <- c(2, 10, 20, 50, 100)
> res.phys.cov(folder, physCovModelFits, gauge, RPs, CI,
  model_type = "Best", locFit = NULL, scaleFit = NULL,
  present_day = TRUE)

# End of complete example script to run all functions

```

References

- Box GE and Cox DR. 1964 'An Analysis of Transformations' Journal of the Royal Statistical Society: Series B (Methodological), 26(2), 211 to 243
- Eastoe EF and Tawn JA. 2009 'Modelling non-stationary extremes with application to surface level ozone' Applied Statistics 58, 22 to 45
- Killick R, Fearnhead P and Eckley IA. 2012 'Optimal Detection of Changepoints With a Linear Computational Cost' Journal of the American Statistical Association, 107(500), 1,590 to 1,598
- Killick R and Eckley IA. 2014 'changepoint: An R package for Changepoint Analysis' Journal of Statistical Software, 58(3), 1 to 19
- Prosdocimi I, Kjeldsen TR and Svensson C. 2014 'Non-stationarity in annual and seasonal series of peak flow and precipitation in the UK' Natural Hazards and Earth System Sciences, 14(5), 1,125 to 1,144
- Rougé C, Ge Y and Cai X. 2013 'Detecting gradual and abrupt changes in hydrological records' Advances in Water Resources, 53, 33 to 44
- Šraj M, Viglione A, Parajka J and Blöschl G. 2016 'The influence of non-stationarity in extreme hydrological events on flood frequency estimation' Journal of Hydrology and Hydromechanics, 64(4), 426 to 437

**Would you like to find out more about us
or about your environment?**

Then call us on

03708 506 506 (Monday to Friday, 8am to 6pm)

email

enquiries@environment-agency.gov.uk

or visit our website

www.gov.uk/environment-agency

incident hotline 0800 807060 (24 hours)

floodline 0345 988 1188 / 0845 988 1188 (24 hours)

Find out about call charges (www.gov.uk/call-charges)



Environment first: Are you viewing this on screen? Please consider the environment and only print if absolutely necessary. If you are reading a paper copy, please don't forget to reuse and recycle if possible.