
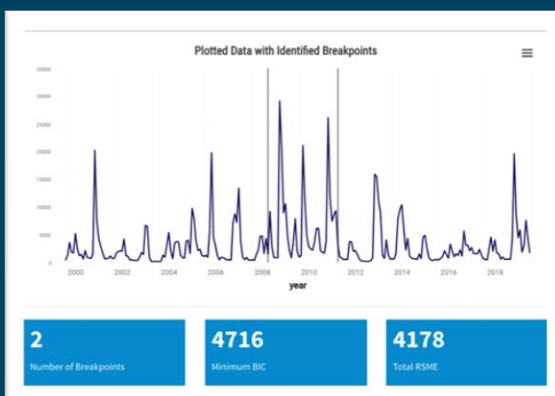





U.S. Army Corps of Engineers

TIME SERIES TOOLBOX USER GUIDE





USER GUIDE

VERSION 2.2, FEBRUARY 2024

U.S. Army Corps of Engineers, Washington, DC
Climate Preparedness and Resilience Community of Practice

Abstract:

The Time Series Toolbox is an online analytic tool for preliminary analysis of time series data, where standard analytic approaches can be applied to user uploaded datasets, preloaded streamflow data, and preloaded stage data. Originally developed to address the need for frequent, but diverse, time series analysis—the tool enables more consistent, repeatable, and efficient time series analysis. Without programming expertise, users can deploy streamlined analysis pipelines, uncovering previously hidden data patterns and rapidly moving from data acquisition to analytic insight. This helps engineers and planners focus on what they do best—applying their expertise to context-centric problems. The tool specifically facilitates the analysis of trend, seasonality, and nonstationarity within time series data. The purpose of this User Manual is to support the use of Version 2.0 of the Time Series Toolbox, released in 2022. The User Manual includes a discussion of the technical concepts incorporated into the Time Series Toolbox and a description of the user interface. This user guide does not cover all possible situations one may encounter using the tool. The first step in conducting time series analysis is to carry out the data upload and preprocessing methods, which are described in Section 1.3. The Time Series Toolbox is not a substitute for deferring to one’s professional engineering judgment.

Preferred Citation:

Olson, Sarah, Nguyen, Marie, Sant-Miller, Aaron, Mueller, Chanel, Veatch, Will & White, Kate (2022)
U.S. Army Corps of Engineers Time Series Toolbox User Guide. U.S. Army Corps of Engineers:
Washington, DC.



TABLE OF CONTENTS

Table of Contents	i
List of Figures	ii
List of Tables	iii
Introduction.....	iv
1 Technical Background	1
1.1 General Capabilities of the Time Series Toolbox	1
1.2 Constraints and Limitations of the Time Series Toolbox	2
1.3 Data Upload and Preprocessing Methods	2
1.4 Trend Analysis	7
1.5 Seasonality	8
1.6 Nonstationarity Detection	9
1.7 Building Time Series Models	21
2 Interacting with the Time Series Toolbox.....	24
2.1 Interactions with Graphs, Data Tables, and Time Filters.....	24
2.2 Explore Data Interface	28
2.3 Model-Based Analysis	31
2.4 Nonstationarity Detection	35
2.5 Time Series Modeling.....	40
3 Example Use Cases	42
3.1 Simple Analysis (with a Priori Knowledge)	42
3.2 Complex Analysis (Without a Priori Knowledge)	47
3.3 No Statistically Significant Nonstationarity Detected	52
Additional References.....	57
Appendix 1: Glossary of Statistics Terms.....	61
Appendix 2: Version Change Log	62



LIST OF FIGURES

Figure 1: Data Entry Panel for User Uploaded Data.....	3
Figure 2: Data Entry Panel for Preloaded Stage and Streamflow Data	4
Figure 3: Selected Site Map in the Data Upload Tab.....	5
Figure 4: Seasonality Preprocessing Options.....	9
Figure 5: Seasonality Components Download.....	9
Figure 6: Data Table Sliders for Filtering Data	25
Figure 7: Typical Data Table Shown Throughout the Application.....	26
Figure 8: Visualization with Export Dropdown Options	27
Figure 9: Timeframe Selection Featured in the Nonstationarity Detection Tab	28
Figure 10: Customization Options in the Upload Panel.....	28
Figure 11: Preprocessing Options in the Data Upload.....	29
Figure 12: Summary Statistics Underneath the Time Series Plot	29
Figure 13: Data Summary Subtab.....	30
Figure 14: Magnificent Seven Information Boxes.....	30
Figure 15: Seasonal Cycle Visualization	31
Figure 16: Trend Analysis Visualization	32
Figure 17: Data Checks for the Trend Analysis Tab.....	33
Figure 18: Decomposition Analysis Options	33
Figure 19: Trend Component of the Seasonality Decomposition	34
Figure 20: Random Component of the Seasonality Decomposition.....	34
Figure 21: Configuration of Parameters for Nonstationarity Tests.....	35
Figure 22: Nonstationarity Detection Graph (First NSD Graph).....	36
Figure 23: Nonstationarity Test Results Graph (Second NSD Graph)	36
Figure 24: Segment Statistics Graph (Third NSD Graph)	37
Figure 25: Gage Drainage Information Box and “View Results and Gage Details” Button on the Nonstationarity Detection Tab.....	37
Figure 26: Additional Gage Details Pop-up on the Nonstationarity Detection Tab	38
Figure 27: Possible Warning Infobox Messages That Appear in the Nonstationarity Detection Tab	38
Figure 28: Method Explorer Individual Test Selector	39
Figure 29: Breakpoint Visualization with Segment Linear Regression Lines Shown	40
Figure 30: Segment Details for Segments in Between Breakpoints	40
Figure 31: Time Series Modeling Diagnostic Plots for Calexico, CA Streamflow Data.....	41
Figure 32: Possible Warning Infobox Messages That Appear in the Nonstationarity Detection Tab	42
Figure 33: Annual Peak Streamflow Time Series for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000).....	43
Figure 34: Nonstationarity Detection Tab for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000)	44
Figure 35: Gage Site Detailed Results for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000)	45
Figure 36: Segment Statistics Graphic for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000)	46
Figure 37: Monotonic Trend Analysis Results for the Full Period of Record 1924-2020.....	47
Figure 38: Annual Peak Streamflow Time Series for the Red River of the North at Grand Forks, ND (USGS ID 05082500)	48
Figure 39: Nonstationarity Detection Tab for the Red River of the North at Grand Forks, ND (USGS ID 05082500)	49



Figure 40: Trend Analysis Tab for the Red River of the North at Grand Forks, ND (USGS ID 05082500):
1882 -2020 51

Figure 41: Trend Analysis Tab for the Red River of the North at Grand Forks, ND (USGS ID 05082500):
Full Period of Record 1882-2020..... 52

Figure 42: Annual Peak Streamflow Time Series for the Upper Iowa River Near Dorchester, Iowa
(USGS ID 05388250) 53

Figure 43: Nonstationarity Detection Tab for the Upper Iowa River Near Dorchester, Iowa (USGS ID
05388250)—Without Timeframe Adjustment 54

Figure 44: Nonstationarity Detection Tab for the Upper Iowa River Near Dorchester, Iowa (USGS ID
05388250)—with Timeframe Adjustment..... 55

Figure 45: Trend Analysis Tab for the Upper Iowa River Near Dorchester, Iowa (USGS ID 05388250) 56

LIST OF TABLES

Table 1: Statistical Methods Implemented in the Nonstationarity Detection Tab 17

Table 2: Monotonic Trend Analysis: Red River of the North at Grand Forks, ND..... 51



INTRODUCTION

Time series data comprises data points measuring the same quantity, indexed in time order. Consequently, time series analysis can be useful to see how a given variable or measurement changes over time. This allows analysts to derive hidden insights and make informed decisions based on patterns in the past data.

Time series data is of particular importance to the United States Army Corps of Engineers (USACE). Many engineering decisions rely on climate, weather, or financial data, which take a sequential structure of reoccurring measurements. The insights found in this data can reveal underlying factors and patterns that impact design and optimization of critical infrastructure.

Climate change is one of many global changes USACE faces in carrying out its missions to help manage the nation's water resources infrastructure, and accordingly, USACE stresses the importance of measuring and studying climate data over time. Water resource managers must make decisions that rely upon assumptions about future supplies, demands, weather, climate, and operational constraints at varying space and time scales. To better enable the required analysis and to empower deeper exploration of time series data more broadly, USACE has developed a suite of tools and web applications for repeatable, quantitative analysis of this data (U.S. Army Corps of Engineers, 2011).

The Time Series Toolbox provides a diverse set of capabilities for time series analysis of both preloaded United States Geological Survey (USGS) streamflow gage site data and user-uploaded data. Primarily, it enables a deeper analysis of trend measurement and seasonality, using regression techniques and statistical methods to identify and define these patterns. The tool also allows for nonstationarity detection by using statistical testing to evaluate change points in the data. This nonstationarity analysis is vital to studying climate data, as there has been increasing evidence that shows that in some places, and for some impacts relevant to the operations of USACE—climate change and human modifications of the watersheds are undermining the fundamental assumption of stationary, resulting in nonstationarity (Milly, et al., 2008). Furthermore, the last section of the tool enables users to visualize and select the most appropriate time series model for the given dataset.

The purpose of this manual is to facilitate the use of Version 2.0 of the Time Series Toolbox, released in 2022. The document includes the following: (1) a discussion of the technical concepts incorporated into the tool, (2) guidance for user interaction, (3) guidance for detecting and interpreting nonstationarities, and (4) a review of example use cases. This user guide does not cover all possible situations that one may encounter while using the tool. The Times Series Toolbox is not a substitute for deferring to one's professional engineering judgment.



ACCESSING THE TOOL

The Time Series Toolbox is available at: https://climate.sec.usace.army.mil/tst_app/

1 TECHNICAL BACKGROUND

As a comprehensive platform, the Time Series Toolbox blends time series analysis, nonstationarity detection, and time series modeling into one analytical tool that improves decision-making within the fields of geological science, hydrology, and engineering. The following sub-sections aim to provide a foundational understanding of the tool, its inputs, and the analysis it enables. The content included in this document should be treated as context required to use the application, not comprehensive technical documentation. Users and readers are encouraged to explore the references included at the end of the document to develop a deeper understanding of the science and engineering theory behind the Time Series Toolbox. In addition, further reading is recommended for users intending to use analytic outputs for future planning or in engineering decision-making.

This technical background section highlights (1) general tool capabilities and limitations, (2) data elements that a user can pass as parameters to the visualizations and data tables, and (3) scientific and engineering background that supports the need for climate scenario projections.

1.1 General Capabilities of the Time Series Toolbox

As a deployed application, the Time Series Toolbox enables the identification and visualization of trends and seasonality within the data. The tool also offers nonstationary analysis, using various statistical tests to detect change points in the user uploaded data. The results of these nonstationarity tests are displayed in three different graphs, which should be analyzed in union to provide a complete analytic picture. The last section of the tool allows users to select one of three time series models to fit and forecast data.

All of these components, working together, can help decision-makers align various climate time series scenarios with existing and planned engineering efforts, estimating when and how the data may impact critical infrastructure and planned development activities.

In practice, the Time Series Toolbox aims to foundationally address three defined needs:

1. Explore the data with a model-based analysis comprising trend detection and seasonality assessment.
2. Create a visualization and summary of nonstationarity analysis of time series data.
3. Aid in the selection of a strong time series model.

To offer data-driven insights, the tool relies on an interactive user interface with four main sections:

1. **Explore Data:** Users can select a preloaded dataset or upload their own dataset (formatted as a CSV file) for which they want to perform analysis in the Explore Data tab. This dataset will be the primary dataset examined throughout the application. Within this panel, there are options for title and axis customization, along with preprocessing options for missing values and data aggregation. Once the data is uploaded, the tool visualizes that data for immediate inspection and evaluation, allowing for further exploration on the Data Summary subtab. This subtab includes a raw data table and more basic summary statistics, including the start date, end date, frequency, missing values, and total



entries. The Magnificent Seven subtab uses Magnificent Seven (Archfield et al., 2013) summary statistics to characterize the uploaded data. Finally, the Seasonal Cycle subtab allows users to explore the data through monthly metric values: mean, minimum, and maximum.

2. **Model-Based Analysis:** This section of the tool contains two subtabs: Trend Analysis and Seasonality. The Trend Analysis section fits trend lines to the data and uses hypothesis testing to measure the significance of the discovered trends in the data. The Seasonality subtab uses a series of statistical methods to detect, extract, and decompose seasonal patterns.
3. **Nonstationarity Detection (NSD):** Within the first subtab, the NSD subtab, the data can be analyzed using nonstationarity detection and testing. The tool analyzes the dataset to see if and when tests detect nonstationarities in the data and then displays what tests detected nonstationarities in which years. The Breakpoint Analysis subtab detects breakpoints in the data using linear regression models and displays the breakpoints as text in a raw data table below an interactive visualization.
4. **Time Series Analysis:** For deeper inspection, the last section helps fit a time series model to the uploaded data. Users can choose between a Time Series Linear Model, an Auto Regressive Integrated Moving Average Model, and an Exponential Smoothing Model, extracting both model fit statistics and the model's forecasts for the uploaded data.

1.2 Constraints and Limitations of the Time Series Toolbox

The Time Series Toolbox facilitates customized time series analysis—offering features to aid users in better understanding their data; however, this tool does not replace engineering decision-making and should be used as a source for the analysis, rather than the sole basis for conclusions.

The tool also has certain requirements and limitations regarding the user inputted data file. Input files are limited to time series of evenly spaced time series data. Furthermore, files with reduced data entries and/or those with several missing values may limit the capabilities of the tool. The tool relies on statistical tests, many of which result in a loss of statistical power when applied to limited datasets. In response, tool functionalities utilizing tests especially sensitive to missing values will either produce an error message or result in interpolation of the data. To leverage the full capabilities of the tool, users are advised to fill in their data with an appropriate interpolation method prior to uploading to the tool.

The Time Series Toolbox is not a substitute for engineering judgment. Engineers are advised to use their judgment to consider the resilience of the system when incorporating the range of results here in their adopted hydrologic study or design results.

1.3 Data Upload and Preprocessing Methods

The Time Series Toolbox facilitates a customized analysis by allowing users to either (a) upload their own data or (b) utilize the preloaded streamflow or stage data. This selection of data source represents Step 1, shown in Figure 1. Immediately after upload, the user can configure the visualization's parameters in the data entry interface. This includes customization for the titles and preprocessing options.

1.3.1 DATA UPLOAD (CUSTOM DATA)

Figure 1 (below) displays the user-uploaded data panel within the Data Upload subtab, which allows users to select a dataset from their local files. The dataset must be in a CSV file format consisting of two columns. These two columns should consist of date and data values, with the date being the first column. The file should include time series data (i.e., a sequence of data points taken at successive equally spaced points in time). The tool currently accepts data of the following frequencies: hourly, daily, weekly, monthly, and yearly.



Users can also select a default dataset to explore the tool and its functionalities. The default dataset used in the application represents discharge data for the Red River of the North at Fargo, ND (USGS 05054000). The data, measured in cubic feet per second, spans a period of monthly data from 2000 to 2019.

Common errors for data formatting include:

1. Data not in CSV format
2. Years not comprising four digits (e.g., year column just containing the last two digits of the year)
3. Missing values in the year column (date is empty, but data column has an entry)
4. Frequency of data is not consistent
5. Data not in sequential order
6. Dataset has more than two columns
7. If hourly data does not work, please ensure data is equally spaced by the hour. If entry still results in error, restructure hourly data to this format: yyyy-mm-ddThh:mm:ss.
8. If hourly data is used, please convert times to Coordinated Universal Time (UTC).
9. Check that the dataset does not have duplicate values (e.g., the same date is entered twice with two different values)

1.3.2 DATA UPLOAD (PRELOADED STAGE AND STREAMFLOW DATA)

Figure 2 (below) displays the upload panel within the Data Upload subtab, which allows users to select a preloaded dataset from USGS sites. The tab allows the user to select the data type and gage location to be used for analysis. The following steps guide the user in selecting a preloaded site.

1. **Select Data Type:** This step is consistent for user uploaded data and preloaded data. For preloaded data, the user has the option to pick between four data types: Annual Instantaneous Peak Streamflow, Annual Stage (Maximum Gage Height), Monthly Mean Stage, and Monthly Mean Streamflow. These datasets are continually updated.

The annual values are updated every 6 months from USGS web-services and stored in the tool. The monthly records are pulled directly from the USGS Water Services API. Therefore, each time a user selects a new monthly dataset, that data will reflect the up-to-date values from USGS.

2. **Select Search Method:** The site selector tab gives the user three methods to search for a site:
 - a. Search by HUC-8: Allows users to filter Site Numbers by HUC-8 alignment
 - b. Search by State: Allows users to filter Site Numbers by State

Figure 1: Data Entry Panel for User Uploaded Data



- c. **Search by Site Number:** Allows users to filter Site Numbers by typing the first couple of characters or the entire Site Number directly

3. **Dynamically Changing Search Method:** Based on the search method selected in (2), the dropdown menu in (3) will dynamically change to render the options corresponding to the selected search method. Once a selection is made in (3), the dropdown menu in (4) will be updated to include all the USGS gage sites with more than 30 years of record in the selected HUC-08 or state. If the selected method in (2) is by Site Number, then (4) will prompt the user to write in the site number. The tabular information populated in the Site Number Information will correspond to the site selected in (4). For example, in Figure 2, the user selected Search by State as the search method in (2). A dropdown to select a state will appear in (3) along with a selection table to select gage in the selected state in (4).
4. **Site Number Information:** The user can view basic metadata associated with the gage site being selected for the analysis (e.g., drainage area, site type). This portion of the site selected tab also provides web links to the selected USGS gage’s “Peak Stream for the Nation” page and the USGS water year summary.
5. **Upload Data:** Once the gage is selected, the user can click the Upload Data button to upload the selected site and view a time series graph of the data. The user can also view the selected site on the “Selected Site Map” (Figure 3). This selected site will be analyzed subsequent tabs.

1. Select Data Source
 Stage (preloaded) = Gage Height, Instantaneous Peak Streamflow (preloaded) = Max Annual Discharge
 Stage (preloaded) ▼

2. Select Search Method
 Search by HUC, State, or Site Number
 Search by State ▼

3. Select State
 Filter gages by state.
 AK ▼

4. Select Site Number
 Note: Not all site numbers are aligned to HUCs or States. Site state alignments are based on USGS data. If desired site number is not found under State or HUC, search by Site Number.
 15052500-MENDENHALL R NR AUKE BAY AK ▼

Upload Data

Site Number Information
[Click here to view more gage details from the USGS website.](#)
[Click here to view the water year summary from USGS website.](#)
 Note: Only gages with available data are shown in the map.

Site Number	Station Name	HUC	State	Country	Drainage Area (sqm)	Site Type
15052500	MENDENHALL R NR AUKE BAY AK	19010301	AK	US	84.8	Stream

Figure 2: Data Entry Panel for Preloaded Stage and Streamflow Data

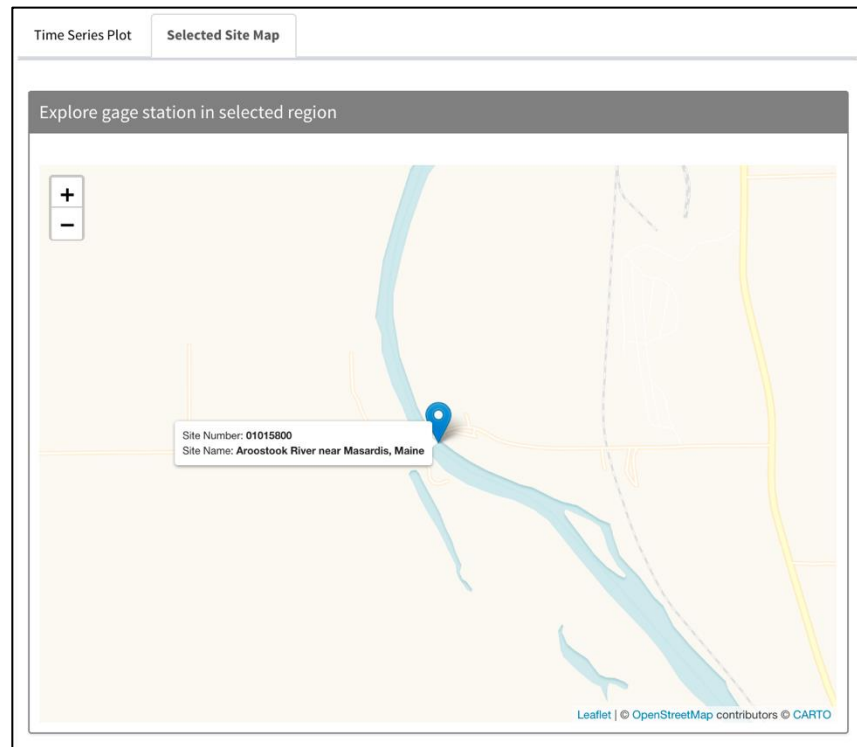


Figure 3: Selected Site Map in the Data Upload Tab

1.3.3 DATA PREPROCESSING

Once the file is uploaded, the user can configure the preprocessing parameters for the Time Series Toolbox. This section of the tool addresses missing values and data aggregation. These preprocessing options are only available for user-uploaded data.

Missing Values

First, users determine how to handle missing values and carry them forward in the analysis. Missing values reflect an incomplete dataset with unknown data values for one or many dates; this lack of information can mislead time series analysis. The options for handling missing data should be considered on a case-by-case basis depending on the dataset and missing values. Different preprocessing options may result in a different analysis.

There are three options the Time Series Toolbox provides to handle missing values:

1. The first option for missing data is to leave the empty entry blank; the analysis will ignore that date and use all other data points for analysis.
2. The second option is imputation, which carries over the previous observation to fill the missing value. This option will result in two or more consecutive years with the same value.
3. The last option is linear interpolation, which uses linear polynomials to construct a new data point based on surrounding data points.

When deciding between which methods to use, users should note the frequency of the data, as well as any trends and variance in the data. For example, if a dataset has many entries and only a couple missing values, then leaving the empty entries blank will not have a huge effect on analysis. However, if these unknown values follow some sort of pattern, then it is possible that not including these values will



introduce bias into the model. Smaller datasets or data with a large number of missing observations may benefit more from linear interpolation or imputation (Grigorios et al., 2018).

Data Aggregation

There are data aggregation options that allow the users to group or roll up their data into larger time bins. These options are only available if the data has a frequency greater than one entry per year. If data aggregation is selected, then the user can select the aggregation level (yearly, monthly, etc.) and then one of the following preprocessing options: cumulative sum, minimum, maximum, or mean.

Detrending

The tool also offers annual detrending for user-uploaded data. For users to access annual detrending, the data needs to be more frequent than one entry per year. Annual detrending works by subtracting the calendar year mean from the individual data points. This option adjusts the data for the annual mean trend.

1.3.4 DATA EXPLORATION

Users can also explore their data in the Data Summary, Magnificent Seven, and Seasonal Cycle subtabs. These subtabs offer various metrics and visualizations for users to better understand and analyze their datasets. Primarily, the Data Summary subtab offers users the ability to examine their data and verify correct entry and aggregation. On the right side of this subtab, information boxes showcase basic summary statistics, including the start date, number of entries, and data frequency. Visualizations of these subtabs are provided in Section 2.2.

The data summary subtab also shows the autocorrelation within the dataset. The autocorrelation statistic used in the tool, provided by the **R stats** package, is defined by the correlation of a signal with a delayed copy of itself. In particular, the lag-1 autocorrelation used in the tool depicts the correlation between data points one time frequency apart (e.g., for weekly data, the correlation between data that is one week apart).

The Time Series Toolbox utilizes the Magnificent Seven summary statistics to help characterize the uploaded data. Research suggests that these measures are both effective and efficient in classifying climate data and generally understanding hydrologic behavior. Below are the seven summary statistics (i.e., calculated as L-moments) and their definitions:

1. **L-Mean:** The average value of the data. Sample L-moments are defined in Hosking (1990). Calculations use the algorithm given in Hosking (1996, p.14).
2. **L-CV (Coefficient of L-Variance):** A measurement of relative variability and is defined as the ratio of the standard deviation to the mean. It takes values from 0-1 and is often expressed as a percentage. If the mean of the data set is zero, the coefficient of variation will approach infinity and hence cannot be calculated. Please see Hosking 1990 and 1996 for more information.
3. **L-Skewness:** The measure of asymmetry of the probability distribution. L-Skewness can be negative, positive, zero, or undefined. A normal distribution has a skew of zero, while a lognormal distribution, for example, would exhibit some degree of right-skew. Please see Hosking 1990 and 1996 for more information.
4. **L-Kurtosis:** The measure of tail density of the probability distribution (i.e., how much of the distribution is contained in the tails). Please see Hosking 1990 and 1996 for more information.
5. **AR1:** The autoregressive lag-one correlation coefficient (i.e., how predictive the previous value in the time series is of the next value). Long-term monthly means are used to remove seasonality from the



data. The code normalizes the data and then applies a first order auto-regression function using the Yule Walker Method.

6. **Amplitude:** A measure of the best fitting, annual sinusoidal curve height. First, flows are standardized, and then fitted to the linear model: $\cos(2\pi t) + \sin(2\pi t)$. The final value is calculated as $\sqrt{\cos(2\pi t)^2 + \sin(2\pi t)^2}$.
7. **Phase:** The measure, in radians, of the angle of the best-fitting, annual sinusoidal curve at time zero. The same pre-processing steps used for calculating Amplitude are used to calculate Phase. However, the final value is calculated as $\arctan(-\sin(2\pi t)/\cos(2\pi t))$.

For more information regarding these statistics, please refer to Archfield et al. (2013), Hosking (1990), and Hosking (1996).

1.4 Trend Analysis

The Time Series Toolbox uses regression techniques to fit monotonic trend lines to the data. This analysis takes part in the Trend Analysis subtab of Model-Based Analysis tab, described in Section 2.3. A monotonic trend implies the dataset consistently increases or consistently decreases; this detected trend may not necessarily be linear. The whole period of record should be evaluated for trends. In addition, detected, strong nonstationarities should be used to subdivide the period of record into stationary subsets, each of which should be tested for the presence of monotonic trends.

Like all statistical methods, these tests and slope calculations are more powerful when applied to larger datasets. None of these methods are appropriate for time series that have fewer than 10 entries. This limitation is because these methods may detect erroneous trends in data with fewer than 10 data points. The significance values employed in applying these tests are a choice among standards in statistics. They are held static (alpha of 0.05) within the tool to facilitate consistency in the application of these tests across USACE. A significance level of 0.05 translates to a 5% probability of encountering a false positive (Type I error): rejecting the null hypothesis given that the null hypothesis is true (e.g., identifying a significant trend when there is no significant trend). In other words, setting the significance level at a certain value predetermines the probability of a Type I error.

1.4.1 TREND SLOPE CALCULATIONS

These trends are calculated with two methods that fit regression curves to the data, effectively measuring trends in the data by extracting the slope coefficient. The tool implements these calculations with the trend package in R (Pohlert, 2015). Values for these two slopes greater than zero correspond to an increasing positive trend, while those less than zero indicate a decreasing negative trend. By using two different methods, the user can lean on a more robust analysis, resulting in higher confidence conclusions.

The two slopes are as follows:

- (1) **Traditional Slope (Least Squares Regression):** Calculated by fitting a simple linear regression to the data and taking the slope of the resulting line. The regression line is fit by minimizing the sum of the squares of the residuals resulting from the proposed equation line and actual data.
- (2) **Sen's Slope:** A robust, nonparametric method to determine the presence of a trend by taking the average of all the slopes between every two points in two-dimensional series. Sen's slope can be significantly more accurate than non-robust traditional slope for skewed and heteroskedastic data.



1.4.2 TREND HYPOTHESIS TESTING

The tool also uses hypothesis testing to measure the significance of the discovered trends in the data. While these tests do not assess the magnitude of the trend, they do evaluate whether the user can confidently say the trend does or does not exist. The tool executes three statistical tests to identify and describe monotonic trends: (1) the Mann-Kendall test, (2) the student's t-test, and (3) the Spearman test. In the Time Series Toolbox, all three of these tests are implemented using the trend package in R (Pohlert, 2015).

Mann-Kendall Test for Trend Detection

The tool implements the nonparametric Mann-Kendall to test for the presence of monotonic trends. The test calculates the Kendall Rank Correlation Coefficient, or Kendall's tau, a value between -1 and 1, where values near -1 suggest a decreasing trend and values close to 1 suggest an increasing trend. The test determines whether the Kendall Rank Correlation Coefficient is significantly different from 0, implying a trend. The alternative hypothesis—that a trend is apparent—is accepted to be true at the 0.05 significance level.

Spearman's Test for Trend Detection

The tool also implements the nonparametric Spearman test to test for the presence of monotonic trends. The Spearman's test calculates the Spearman Rank Correlation Coefficient, or Spearman's rho, a value between -1 and 1. Values close to -1 suggest a negative trend and values close to 1 suggest a positive dependence between two variables. Like the Mann-Kendall test, this test is applied to determine whether the Spearman Rank Correlation Coefficient is significantly different from 0, implying a trend. The alternative hypothesis—that a trend is apparent—is accepted to be true at the 0.05 significance level.

Student T-Test for Trend Detection

In addition to implementing the nonparametric Mann-Kendall and Spearman tests, the Student's t-test is applied to evaluate the datasets for linear trends. The Student's t-test is parametric and relies on the assumption of normality. The t-test is applied to evaluate the validity of fitting a linear regression relationship to the data with a given slope. The t-test evaluates the null hypothesis that the slope can be defined as zero, which would imply that no trend exists in the data (the independent and dependent variables being assessed are uncorrelated). Under the Trend Analysis tab, whether the linear regression slope is negative or positive is indicative of the directionality of any statistically significant trend. The alternative hypothesis—that a trend is apparent—is accepted to be true at the 0.05 significance level.

1.5 Seasonality

The Time Series Toolbox uses a series of statistical methods to identify and define seasonal patterns. These techniques consider underlying trends in the data, as well as noise and natural volatility. In particular, the tool uses an additive decomposition model that separates out three of the different components embedded within time series data, where:

$$Data \sim Trend + Seasonality + Error$$

These three components are defined as:

- (1) **Trend:** Consistent increases/decreases in the data as determined by a filter or smoother, such as a moving average



(2) **Seasonality:** Recurring similarity of the data between time periods (e.g., months) after the removing the trend

(3) **Random:** Any increases/decreases in the data after the trend and seasonality is removed

The tool uses two different methods to decompose the original data to better capture seasonality, trends, and random effects. Using both methods in one analysis allows for a more robust analysis of the data and conclusions with higher analytical confidence. The two methods are: (1) Moving Average Decomposition and (2) Seasonal-Trend Decomposition by Local Polynomial Regression (STL).

1. **Moving Average Decomposition:** The moving averages (a succession of averages derived from successive segments) are fit to the data to identify both trends and seasonality.
2. **STL Decomposition:** Curves are fit to the data to identify both trends and seasonality. The STL method is often more robust to outliers, so that occasional unusual observations will not affect the estimates of the trend-cycle and seasonal components (Cleveland, R. B. et al., 1990).

For more information regarding seasonal decomposition, please refer to Hyndman, R. J., & Athanasopoulos, G. (2018). In particular, Chapter 6 discusses time series decomposition, the Moving Average method, and the STL method.

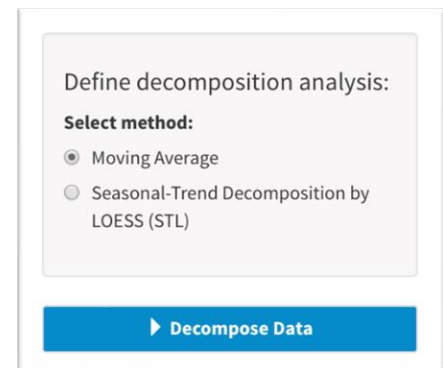


Figure 4: Seasonality Preprocessing Options

After the signal is seasonally decomposed, the user can download the results of the decomposition in a csv file. The csv file is formatted with columns for the date, original data, trend component, seasonality component, and random component. Additionally, the time series plots for each of the seasonally components can be downloaded as one image.

If the user wants to run time series analyses on one of the seasonally decomposed components, the user can download the seasonally decomposed signal and reupload it on the Explore Data tab. Alternatively, the user can select the “Rerun Analysis” button as seen in Figure 5. Once that button is clicked, the user will be navigated to the Explore Data tab with the component selected preloaded. Along with the three seasonally decomposed components (trend, seasonality, and random), the user can also select “Seasonality Removed”, which contains the data without seasonality (original data minus the seasonality component).

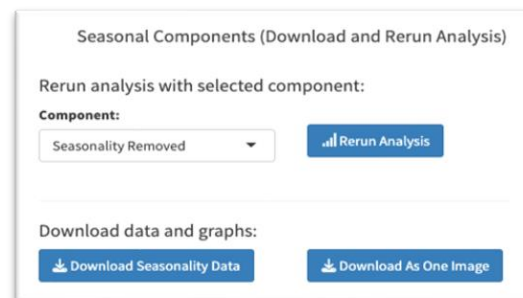


Figure 5: Seasonal Components Download

1.6 Nonstationarity Detection

The Time Series Toolbox uses statistical testing to detect the presence of nonstationarities in the uploaded data. Although the stationarity assumption—which implies the statistical characteristics of time series data are constant throughout time—has been foundational to climate and engineering decision-making in the past, recent scientific evidence shows that in some places, climate change and human modifications are undermining this fundamental assumption, resulting in nonstationarity in the time series (Milly et al., 2008). This section of the Time Series Toolbox allows users to detect if there are any nonstationarities



present in the data, which can significantly impact decision-making about future series behavior. The detection is done using statistical tests that evaluate for changes in the mean, variance, or distribution in the data.

The tool performs nonstationarity detection using the open-source R statistical programming language (R Development Core Team, 2011) and four freely available statistical test packages associated with it: *cpm* (Ross, 2015), *bcp* (Erdman & Emerson, 2007), *ecp* (James & Matteson, 2015), and *trend* (Pohlert, 2015). These packages contain the following sets of statistical methods that the tool uses to detect nonstationarities in streamflow and stage records: likelihood ratio testing, energy-based algorithms, Bayesian Monte-Carlo simulation, and Pettitt tests. In addition, a custom package developed specifically for this tool executes the Lombard model to test for abrupt and smooth nonstationarities.

1.6.1 THE STATIONARY ASSUMPTION

A stationary process assumes that the mean, variance, and autocorrelation structure do not change over time. A stationary dataset can be relatively easy to predict and analyze, as evaluations of future behavior can be predicated on historic patterns. However, evidence shows that some datasets cannot be analyzed under the stationary assumption; instead, the data should be analyzed as nonstationary. As many time series techniques assume stationary behavior, the detection of nonstationarities in a series is a critical element of the analytic workflow (Apurv, T., & Cai, X., 2019).

1.6.2 NONSTATIONARITY IN CLIMATE AND ENGINEERING ENVIRONMENTS

Recent scientific evidence shows that substantial anthropogenic change of Earth's climate is altering the means and extremes of temperature, precipitation, evapotranspiration, and rates of discharge of rivers. These impacts mirror those previously seen in climate data, as the stationarity assumption has long been compromised by human disturbances in river basins. Flood risk, water supply, and water quality are affected by water infrastructure, channel modifications, drainage works, and land-cover and land-use change (Milly et al., 2008).

Understanding whether a dataset includes nonstationary behavior is vital for time series analysis, as the subsequent research and engineering decisions depend on these results. Coastal communities depend on accurate reporting and prediction measures. For example, studies have shown that the assumption of stationarity can lead to underestimation of drought risk, thereby exposing water resource systems to failure under severe droughts (Apurv, T., & Cai, X., 2019).

While the Time Series Toolbox currently only provides preloaded streamflow and gage height datasets, a literature review was performed to confirm the validity of applying nonstationarity detection on meteorological variables, including temperature and precipitation. There are many examples of the scientific community applying nonstationarity detection tests to various temperature and precipitation variables, at a variety of temporal resolutions (e.g., Kumar et al., 2022, Yaman, B., & Ertuğrul, M., 2020, and Zhou et al., 2022). While temperature and precipitation datasets are not currently preloaded into the tool, the user can load these data types into the tool to perform the analyses available so long as the dataset follows the requirements for a time series.

1.6.3 TEMPORAL RESOLUTION CONSIDERATIONS

As nonstationarity detection tests are sensitive to temporal resolution, users should consider how the temporal resolution of their datasets may impact the performance of nonstationarity detection. Sub-annual time series (e.g., monthly) can contain a seasonal component. A literature review of nonstationarity detection tests run on time series at different temporal resolutions revealed that seasonality can influence tests' ability to detect nonstationarities. Lund et al. (2007) shows the reduced power (probability that a



test correctly rejects the null hypothesis) of nonstationarity detection tests when applied to simulated, periodic time series with added changepoints as opposed to simulated, non-periodic time series with added changepoints. The periodicity added to simulated time series is meant to represent seasonality in a dataset.

Due to the negative effects of seasonality on the accuracy of nonstationarity detection, seasonality is typically removed from a time series before performing nonstationarity detection. Based on a literature review, there are two methods commonly used for handling seasonality: (1) model and remove seasonality from the signal before performing nonstationarity detection (Lund et al., 2007 and Militino et al., 2020), or (2) split data into per month (e.g., January time series) or per season (e.g., Winter time series) time series and perform nonstationarity detection on each time series individually (e.g., Kumar et al., 2022 and Yaman, B., & Ertuğrul, M., 2020). The second option involves splitting a sub-annual time series into multiple time series. For example, a monthly time series could be split into 12 separate time series (one per month) where each time series would contain the yearly values for a given month (e.g., January time series contains data from each January within a year range). The nonstationarity detection tests would then be run on each of the 12 per-month time series separately.

Another way to remove seasonality in the data is to perform nonstationarity detection on data aggregated to the annual scale. However, an analysis of nonstationarity detection results on per-month, per-season, and annual time series showed that annual-resolution timeseries may mask nonstationarities at the sub-annual scale (i.e. monthly or seasonal scale). While this is not necessarily an issue, users of the Time Series Toolbox should consider what temporal scale makes the most sense for their project and take into account the effects of seasonality in the case of sub-annual time series.

1.6.4 TYPES OF STATISTICAL METHODS FOR NONSTATIONARITY DETECTION

The statistical methods applied within the Time Series Toolbox are categorized by the type of nonstationarity detected and whether the nonstationarity is representative of a smooth or abrupt change in the statistical properties of the dataset under analysis.

1. **Type of Nonstationarity:** Methods that only detect differences in the mean will not detect nonstationarities where the mean of the data stays consistent while the variance/standard deviation changes. Likewise, methods that only detect differences in variance/standard deviation will not detect nonstationarities where the mean of the data changes while the variance/standard deviation stays consistent. For this reason, the tool applies several methods that assess nonstationarities in time series datasets driven by changes in the mean, variance/standard deviation, and in the overall, distributional properties of the dataset.
 - a. **Mean-Based Nonstationarity Methods** detect a nonstationarity where there has been a significant change in the average value of the data. The Lombard Wilcoxon, Pettitt, Mann-Whitney, and Bayesian methods are all mean-based.
 - b. **Variance-Based Nonstationarity Methods** detect a nonstationarity where there has been a significant change in the variance/standard deviation of the data. The Mood and Lombard Mood methods are variance-based statistical tests.
 - c. **Distribution-Based Nonstationarity Methods** detect a nonstationarity where there has been a significant change in the underlying distribution of the data, which can be caused by a change in the parameters of the same distribution (e.g., going from one normal distribution to another) or from one distribution type to another (e.g., going from a normal distribution to a beta distribution). The Cramer-Von-Mises, Kolmogorov-Smirnov, LePage, and Energy Divisive methods are distribution-based.



2. **Smooth versus Abrupt Nonstationarities:** The methods are also categorized by whether the nonstationarities they detect represent abrupt or smooth changes in the data. Here, a smooth transition refers to a gradual change in the mean, variance/standard deviation, and/or distribution within the datasets. The only methods in the tool that can detect smooth changes in the statistical properties of the datasets under analysis are the smooth Lombard Wilcoxon and Smooth Lombard Mood methods. The remaining methods (including the abrupt Lombard Wilcoxon and abrupt Lombard Mood methods) detect abrupt changes in the dataset. Abrupt changes occur at a single point in the record and separate the period of record into two subsets of data that have one or more of the following: different means, different variances and standard deviations, or different distributions. A record can have multiple abrupt changes, separating it into several subsets of data.

1.6.5 OTHER CHARACTERISTICS OF STATISTICAL METHODS FOR NONSTATIONARITY DETECTION

In addition to the categorization of statistical methods described above, there are other defining characteristics associated with the statistical tests applied in the Time Series Toolbox. These traits should be taken into consideration when interpreting the results generated by the tool, especially when assessing why some methods detected a nonstationarity while other methods did not. The properties associated with each of the statistical tests packaged within the Nonstationarity Detection tab are described in further detail in Table 1.

1. Methods are either parametric or nonparametric. Parametric statistical tests rely on the assumption that a dataset fits a known statistical distribution (i.e., normal distribution). If the data used for detection does not meet this assumption, then parametric methods should not be applied because results might be inconsistent or incorrect. Hydrometeorological time series datasets are typically difficult to characterize using a statistical distribution. For this reason, the majority of the nonstationarity detection tests applied by the tool are nonparametric. The one exception is the Bayesian method, which relies on the assumption of normality. Within the Time Series Toolbox, the Anderson-Darling Test is applied to assess whether a given dataset conforms to a normal distribution. The tool will not apply the Bayesian method, unless the results of the Anderson-Darling Test indicates that the selected record conforms to a normal distribution.
3. Methods can detect either single or multiple nonstationarities. The Pettitt method and all of the Lombard methods are only able to detect the single most statistically significant nonstationarity in a time series. The rest of the methods can detect multiple nonstationarities within a single time series.
4. Statistical methods applied to identify abrupt change points in the statistical properties of the datasets being analyzed can indicate a detected change point based on the year prior to (end year) or after (first year) the change in mean, standard deviation/variance, or statistical distribution. If a method identifies the shift in a dataset's statistical properties by indicating the first year after the change in mean, standard deviation/variance or distribution, then "Yes" appears in the "Detects First Year" column of Table 1. For guidance on how to account for this difference across methods, refer to the matrix in Section 1.6.9.
5. Statistical tests differ in sensitivity, specificity, and assumptions. Because every statistical test included in the Time Series Toolbox has its own strengths and weaknesses, no individual test can be counted on for every situation. For this reason, an array of 12 different statistical tests is incorporated into the tool to account for the shortcomings of any one individual test.

1.6.6 STATISTICAL SIGNIFICANCE

The Nonstationarity Detection tab is designed to identify statistically significant nonstationarities within the data. The way we measure this statistical significance differs with each nonstationarity detection method. The parameters can be adjusted using the Sensitivity Parameters within the NSD tab, further



described in Section 2.4.1. For more details on these differences and how adjustments to sensitivity parameters affect the output, please refer to Table 1 in Section 1.6.10.

1. As is typical for a Bayesian approach, the Bayesian change point (**bcp**) method assumes a preliminary, also known as a “prior,” statistical distribution. Then, the method iteratively samples data to adjust this “prior” distribution into a “posterior” adopted distribution. By repeatedly partitioning the full time series into subsets and carrying out the Bayesian approach to define the “posterior,” for each subset of data, the test can evaluate the statistical properties of each of the segments of data. If the partitions have “posterior” distributions with different means, the bcp flags their boundaries as nonstationarities. Within the tool, two parameters can be adjusted to alter the sensitivity of the bcp method.

The first of these two parameters is the “Bayesian Posterior Threshold” parameter (Erdman & Emerson, 2007), which acts as a threshold defining the test’s sensitivity for nonstationarity detection. This parameter represents the minimum likelihood needed to determine the presence of a nonstationarity based on the “posterior” distribution. A default level of 0.5 is recommended. A value of 0.5 implies neutrality, meaning it does not favor the presence or absence of a nonstationarity. Higher or lower values would represent favoring the absence (for lower values) or presence (for higher values) of nonstationarities. Note that the “Bayesian Posterior Threshold” parameter does not impact underlying “posterior” distribution itself, but rather represents a threshold for significance.

Users can also choose to adjust the “Bayesian Prior Likelihood” parameter. This parameter refers to the prior likelihood of an uncertain quantity: a probability distribution reflective of deductive reasoning before empirical evidence is considered (Erdman & Emerson, 2007). A prior of long duration will make it hard to detect change points, while one of short duration (the distance between successive change points) may produce spurious change points. Studies from Yao 1984 and Barry/Hartigan 1993 recommend setting the Bayesian Prior Likelihood parameter to 0.2, which is the default value in the Time Series Toolbox (Erdman & Emerson, 2007). However, users may choose an alternate sensitivity parameter value based on professional judgment and experience.

2. Both the smooth and the abrupt Lombard methods carry out a hypothesis test to evaluate the validity of the assumption that the time series being analyzed contains a nonstationarity. Lombard Wilcoxon checks for a change in mean and Lombard Mood checks for a change in variance/standard deviation (Quessy, Favre, Said, & Champagne, 2011). If the Lombard hypothesis test returns a true result at the selected level of significance, the Lombard method will return a detected nonstationarity.

The selected level of significance for the Lombard method represents the probability of a Type I error. A Type I error is the probability of accepting an alternative hypothesis when the null hypothesis is true. In this context, it is the probability of incorrectly asserting that the record’s mean or variance/standard deviation is changing when in fact it is stationary. Because the Lombard method sensitivity parameter represents the acceptable probability of a Type I error, 0.05 is adopted as the tool’s default to be consistent with the significance level typically adopted within statistical literature. The user can change the default value to 0.01 or 0.1 based on professional judgment and experience.

3. The energy divisive method (**ecp**) tests for nonstationarity by iteratively dividing the dataset being analyzed into a series of segments. To determine if a nonstationarity is detected at a given point, the method repeatedly performs a statistical test against all previous segments to determine if a reasonable level of statistical significance has been achieved. After identifying a statistically significant nonstationarity, the portion of the time series following the nonstationarity is evaluated for



subsequent nonstationarities. The process continues until it reaches a point that is not statistically significant (James & Matteson, 2015).

The **ecp** level of statistical significance represents the probability of a Type I error, or the probability of detecting a nonstationarity when none is present. To account for the repeated tests across all possible subsets of the period, a higher level of default sensitivity (0.5) has been chosen to ensure that the method regularly returns at least one nonstationarity. The user can rely on professional judgment to adjust the sensitivity parameter.

4. For the Pettitt test, the tool returns a single nonstationarity at the point in the dataset where the difference in the mean between the subsets of data prior to and after the nonstationarity has the greatest level of statistical significance (or smallest probability of a Type I error). The significance value must also be below the selected sensitivity threshold for statistical significance. Type I errors represent the probability of accepting an alternative hypothesis when the null hypothesis is true. In this context, it is the probability of incorrectly asserting that there is a nonstationarity when in fact there is none. Because the Pettitt method sensitivity parameter represents the acceptable probability of a Type I error, 0.05 is adopted as the tool's default to be consistent with the significance level typically adopted within statistical literature. The user can change the default value to 0.01 or 0.1 based on professional judgment and experience.
5. For the Change Point Model (CPM) methods, the algorithm runs the named statistical test (e.g., the Mood test) on the two segments of the time series created by partitioning the record at every possible point in the series to initially detect changes in the mean, variance/standard deviation, or distribution. After changes are detected, the data is re-examined to determine when these detected changes first occurred in sequence. This process only returns statistically significant results based on the sensitivity parameter selected by the user. This sensitivity parameter corresponds to the average number of data points until a Type I error occurs (i.e., average run length), or the probability of accepting an alternative hypothesis when the null hypothesis is true (Ross, 2015). The recommended default value is 1,000, which corresponds to an alpha of 0.001. However, the user can change this sensitivity parameter to be one of the following values: 370, 500, 600, 700, 800, 900, 1,000, 2,000, 3,000, 4,000, 5,000, 6,000, 7,000, 8,000, 9,000, 10,000, 20,000, 30,000, 40,000, or 50,000. The values for this sensitivity parameter are limited as they are related to a series of precomputed statistical distributions. The choice of an alternate sensitivity parameter may be guided by professional judgment.

If the user has prior knowledge of a nonstationarity in the data (e.g., construction and operation of a dam with significant storage upstream from a stream gage), then that expected nonstationarity could be used as a guide for determining the associated level of significance applied by each of the statistical tests described above. The default levels of significance described above may not be able to detect the known nonstationarity. In this case, the user is advised to adjust the sensitivity parameters associated with each test using professional judgment and experience to ensure they are sensitive enough to detect this known nonstationarity. After adjustments to sensitivity parameters are made, any additional nonstationarities detected could be considered equally statistically significant.

The detection of multiple nonstationarities within a short span is partly due to differences between how each method detects and identifies nonstationarity years. Some methods detect the last year of a homogenous period, while others identify the first year of a subsequent homogenous period (Table 1). Where the various abrupt statistical methods detect statistically significant nonstationarities within five years of one another—or if the smooth Lombard method detects a transition period over which the statistical properties of the dataset are changing that is less than five years long—the user is encouraged to



“combine” the detected nonstationarities to be representative of a single change point in the record. To select which year to use as the representative nonstationarity, consider the following criteria:

- If more methods detected a nonstationarity in one year than in the other, then the year with more concurrence is best suited to segment the data.
- If multiple years have equal concurrence among the methods (e.g., two methods detected nonstationarities in each of two consecutive years), then the more recent year is a more appropriate choice to segment the data.

Similarly, when the smooth nonstationarity detection techniques identify a statistical transition period of five years or less, this can be considered indicative of a single nonstationarity, and the change point should be selected as the last year being identified as part of the smooth transition.

At some sites, the tool may identify multiple clusters of nonstationarity years and/or long periods of time over which the smooth nonstationarity detection techniques indicate statistical instability. Statistically significant results are not necessarily operationally significant. The magnitude of the shift in the mean and/or variance/standard deviation should be considered relative to study purpose and/or level of risk adverseness when evaluating whether a nonstationarity is operationally significant.

1.6.7 LENGTH AND CONTINUITY OF PRELOADED USGS RECORDS

The preloaded USGS streamflow datasets available within the tool include annual instantaneous peak streamflow, annual maximum stage, monthly mean stage, and monthly mean streamflow. Although the tool pulls all sites with more than three years of data, it is not recommended to run nonstationarity analysis on data with less than 30 years of record. Records being adopted for analysis should also ideally contain no discontinuities in the peak streamflow or stage record. If the period of record contains significant discontinuities (more than 2-3 consecutive years OR missing >5% of the record), or if it has been truncated by the user to less than 30 years of data, then detected nonstationarities may be inaccurate. The tool automatically warns the user if there are gaps in the USGS annual instantaneous peak data record at a given gage site or when statistical tests for nonstationarities are being applied to less than 30 years of record.

The statistical methods used in the tool require a continuous record to detect changes in the statistical properties of the dataset. If missing data is detected, the tool will omit the missing values from the period of record and analyze the remaining data as if it was a single continuous sequence. For example, if a period of record is complete from 1902 to 2002, but the peak streamflow data for 1967 is missing, the tool will shift all observations after 1968 one year earlier to create a single continuous sequence for nonstationarity detection.

In addition to the sensitivity parameters described above, the CPM methods have an additional “CPM Methods Burn-In Period” parameter. The default value for this parameter is 20. For example, with a burn in period of 20, any changes in data series will not be detected until after the twentieth observation in the time series. The default value of 20 provides a good balance between (a) providing enough time for statistical methods to incorporate the variability in the record and (b) not missing any early nonstationarities. However, if users would like to detect change points earlier than this period, they can adjust this value between 5 and 50 using the “CPM Methods Burn-In Period” slider on the Sensitivity Parameters section of the Nonstationarity Detection tab.



1.6.8 METHODS FOR NONSTATIONARITY DETECTION

The Time Series Toolbox uses several tests to evaluate the mean, median, scale, or distribution of the data over time. Below are the tests used to detect nonstationarities in these measures:

Mean-Based Tests

1. **Bayesian Change Point Detection (CPD):** Parametric (Gaussian) test that uses product partitions to identify change points within a sequence using MCMC sampling by assuming a sequence can be broken into partitions with a constant mean, where changes in the mean between partitions are change points (Erdman & Emerson, 2007). For the Bayesian CPD test, we use the **bcp** package in R.

Median-Based Tests

1. **Lombard Wilcoxon:** Test that nests the Wilcoxon score function within the Lombard test statistic to detect both smooth and abrupt shifts in median by time. It determines the level of significance at which nonstationarities should be detected. It can also be interpreted as the maximum p-value associated with these nonstationarities (Quessy, J.-F. et al., 2011). For the Lombard Wilcoxon test, we use a custom script in R.
2. **Pettitt:** Nonparametric test that identifies change point in the median by testing whether two samples come from the same population. The tool returns a single nonstationarity at the point in the dataset where the difference in the median between the subsets of data prior to and after the nonstationarity has the greatest level of statistical significance or smallest probability of a Type I error (Tan, X. and T.Y. Gan, 2015). For the Pettitt test, we use the **trend** package in R.
3. **Mann-Whitney:** Nonparametric test on median shift that tests if a randomly selected value from one sample is greater than or less than a randomly selected value from the comparison sample (seen as the nonparametric counterpart to the t-test) (Ross, 2015). For the Mann-Whitney test, we use the **cpm** package in R.

Scale-Based Tests

1. **Mood:** A nonparametric case of a Pearson's Chi-test that evaluates change points based on volatility in medians between defined samples (Ross, 2015). For the Mood test, we use the **cpm** package in R.
2. **Lombard Mood:** Nests the Mood score function within the Lombard test statistic to detect both smooth and abrupt shifts in scale by time (Quessy, Favre, Said, & Champagne, 2011). For the Lombard Mood test, we use a custom script in R.

Distribution-Based Tests

1. **Cramer-von-Mises:** Nonparametric goodness of fit test that compares two empirical distributions by evaluating a test statistic of distributional distance (Ross, 2015). For the Cramer-von-Mises test, we use the **cpm** package in R.
2. **Kolmogorov-Smirnov:** Nonparametric test that compares two empirical distributions by evaluating a test statistic of distributional distance (Ross, 2015). For the Kolmogorov-Smirnov test, we use the **cpm** package in R.
3. **LePage:** Simultaneously tests the equality of both the location and scale parameters, where an inequality in one suggests distributional shift (Ross, 2015). For the LePage test, we use the **cpm** package in R.
4. **Energy Divisive:** Nonparametric test based on hierarchical clustering, where change points are iteratively identified and can be diagrammed as a binary tree. The statistical significance is examined by means of a permutation test that combines bisection and multivariate divergence measures (James & Matteson, 2015). For the Energy Divisive test, we use the **ecp** package in R.



Smooth-Based Tests

These testing methods are also categorized by whether the nonstationarities they detect represent abrupt or smooth changes in the data. A smooth transition refers to a gradual change in the mean, variance/standard deviation, and/or distribution within the record. The only methods in the tool that can detect smooth changes in the statistical properties of the datasets under analysis are the smooth Lombard Wilcoxon and smooth Lombard Mood methods. The remaining methods (including the abrupt Lombard Wilcoxon and abrupt Lombard Mood methods) detect abrupt changes in the dataset. Abrupt changes occur at a single point in the record and separate the period of record into two subsets of data that have one or more of the following: different means, different variances and standard deviations, or different distributions. A record can have multiple abrupt changes, separating it into several subsets of data (Friedman, D. et al., 2018).

1.6.9 METHODS FOR NONSTATIONARITY DETECTION MATRIX

Table 1 provides an overview of each of the statistical methods applied by the Time Series Toolbox for nonstationarity detection. Please note that those methods that share names with established statistical tests are algorithms that apply the named statistical test repeatedly to detect nonstationarities. The algorithms include a check for whether the p-values from their respective tests demonstrate statistically significant results, as detailed in Section 1.6.6.

Table 1: Statistical Methods Implemented in the Nonstationarity Detection Tab

Statistical Test	Parametric	Type of Nonstationarity	Single or Multiple Nonstationarities	Detected Year	Input Parameters			Strengths/Weaknesses
				Last = Last year prior to detected shift First = First year after detected shift	Parameter	Default	Description	
The Change Point Model (CPM)								
Cramer-von-Mises	No	Distributional/ Abrupt	Multiple	Last	Burn-In Period	20	Determines how many observations are used for the initialization period. If startup = 20, then changes that occur in the first 20 observations will not be used for change point detection. As this parameter is specifically related to the initialization period, the final change points detected by this method may occur prior the value supplied by the user. This is used because the statistical tests will have low reliability for small sample sizes.	This method may be more likely to detect a nonstationarity related to the tails or extremes of a distribution rather than the center of a distribution.



Statistical Test	Parametric	Type of Nonstationarity	Single or Multiple Nonstationarities	Detected Year	Input Parameters			Strengths/Weaknesses
				Last = Last year prior to detected shift First = First year after detected shift	Parameter	Default	Description	
					CPM Sensitivity	1000	Higher values will result in fewer nonstationarities. Must have one of the following values: 370, 500, 600, 700, 800, 900, 1,000, 2,000, 3,000, 4,000, 5,000, 6,000, 7,000, 8,000, 9,000, 10,000, 20,000, 30,000, 40,000 or 50,000.	
Kolmogorov-Smirnov	No	Distributional/Abrupt	Multiple	Last	See CPM Cramer-von-Mises			This method may be more likely to detect a nonstationarity related to the center of a distribution rather than the tails or extremes of a distribution.
Mann-Whitney	No	Median/Abrupt	Multiple	Last	See CPM Cramer-von-Mises			
LePage	No	Distributional/Abrupt	Multiple	Last	See CPM Cramer-von-Mises			
Mood	No	Scale/Abrupt	Multiple	Last	See CPM Cramer-von-Mises			
The Lombard Model								
Lombard Wilcoxon	No	Median/Abrupt	Single	Last	Lombard Sensitivity	0.05	This determines the level of significance at which nonstationarities should be detected. It can also be interpreted as the maximum p-value associated with these nonstationarities.	This method does not work well with short time series, or with small changes in magnitude.
		Median/Smooth						
	No	Scale/Abrupt	Single	Last	See Lombard Wilcoxon			
Lombard Mood		Scale/Smooth						This method does not work well with short time series, or with small changes in magnitude.
Other Methods								



Statistical Test	Parametric	Type of Nonstationarity	Single or Multiple Nonstationarities	Detected Year	Input Parameters			Strengths/Weaknesses
				Last = Last year prior to detected shift First = First year after detected shift	Parameter	Default	Description	
Pettitt	No	Median/Abrupt	Single	First	Pettitt Sensitivity	0.05	This represents the probability that the nonstationarity detected is a false positive or the p-value associated with the Pettitt test.	
Bayesian Change Point (BCP)	Yes	Mean/Abrupt	Multiple	Last	Bayesian Posterior Threshold	0.5	This is the minimum likelihood needed to determine the presence of a nonstationarity, based on the "posterior" distribution. A value of 0.5 implies neutrality, meaning it does not favor the presence or absence of a nonstationarity.	This method does not work well with short time series, or with small changes in magnitude. This method can only be applied to datasets which fit a normal distribution.
					Bayesian Prior Likelihood	0.2	The prior likelihood represents our prior beliefs about a value—without any empirical evidence—modelled into a probability distribution (Erdman & Emerson, 2007). As recommended by other studies, the default value is set to 0.2 (Yao 1984, Barry/Hartigan 1993).	
Energy-Based Divisive Method (ECP)	No	Distributional/Abrupt	Multiple	First	Energy-Based Sensitivity	0.5	This value represents a change threshold or the probability that the point in question is a nonstationarity. Larger values will yield less nonstationarities.	This method can sometimes detect a series of nonstationarities located in close proximity to one another. These nonstationarities may not all be operationally significant, so it is important to examine them carefully.

1.6.10 TUNING DETECTION METHODS

In addition, these tests can be modified by changing the parameters. Some of the parameter changes will result in fewer nonstationarities, while other parameter changes may result in more nonstationarities detected. For the following parameters, larger values will decrease the sensitivity of these tests, resulting in fewer nonstationarities detected:



1. **CPM Methods Burn-In Period:** If CPM Methods Burn-In Period is set to 20, this means that any changes in the data series cannot be detected during the initialization period by these methods until after the twentieth observation in the time series. However, the resulting nonstationarities identified by this method may occur prior to this value.
2. **CPM Methods Sensitivity:** This sensitivity parameter corresponds to the average number of data points until a Type I error occurs, or the probability of accepting an alternative hypothesis when the null hypothesis is true. The values for this sensitivity parameter are limited as they are related to a series of precomputed statistical distributions (Ross, 2015).
3. **Bayesian Posterior Threshold:** This parameter does not affect the underlying assumptions in the method, but it does act as a threshold for nonstationarity detection, or the minimum likelihood needed to determine the presence of a nonstationarity based on the “posterior” distribution (Erdman & Emerson, 2007).
4. **Energy Divisive Method Sensitivity:** Similar to the CPM sensitivity parameter, a sensitivity parameter has been implemented to adjust the level of statistical significance this method uses. To account for the repeated tests across all possible subsets of the period, a higher level of default sensitivity (0.5) has been chosen to ensure that the method regularly returns at least one nonstationarity. The user can—using professional judgment—adjust the level of statistical significance this method uses by adjusting the accompanying sensitivity parameter (James & Matteson, 2015).

For these parameters, smaller values will decrease the sensitivity of these tests, resulting in fewer nonstationarities detected.

1. **Pettitt Sensitivity:** For the Pettitt test, the tool returns a single nonstationarity at the point in the dataset where the difference in the mean between the subsets of data prior to and after the nonstationarity has the greatest level of statistical significance (or smallest probability of a Type I error). The significance value must also be below the threshold set in the Pettitt sensitivity parameter (Friedman et al., 2016).
2. **Bayesian Prior Likelihood:** The prior likelihood of an uncertain quantity is the probability distribution that would express one's beliefs about this quantity before some evidence is taken into account (Erdman & Emerson, 2007).
3. **Lombard Smooth Methods Sensitivity:** If the hypothesis test returns a true result at the selected level of significance, the Lombard method will return the detected nonstationarity. The selected level of significance for the Lombard methods represents the probability of a Type I error, or the probability of accepting an alternative hypothesis when the null hypothesis is true. In this context, it is the probability of incorrectly asserting that the record contains nonstationarities, when in fact, it is stationary. The default value for the Lombard method sensitivity in the tool is 0.05, but the user can change this to 0.01 or 0.1.

For more information about these tests and parameter tuning, please refer to Friedman et al., 2016.

1.6.11 APPLYING FINDINGS

When making engineering decisions, nonstationarity detection is now a critical phase in the analytic process. If nonstationarities are detected, one must be cautious in concluding any historical analysis is representative of future behavior. It is up to the user to determine which, if any, of the statistically significant nonstationarities identified by the Time Series Toolbox may be used to segment the data for analysis. It is advisable to assess the relative “strength” of any nonstationarities detected to determine which should be factored into subsequent decision-making.



A “strong” nonstationarity is one for which there is a consensus among multiple nonstationarity detection methods, robustness in detection of changes in statistical properties, and relatively large change in the magnitude of a dataset’s statistical properties. As there are many types of nonstationarity, guidance recommends the use of many tests at one time, as each test has its own unique tradeoffs. When testing for nonstationarity, researchers generally use three identification criteria:

1. **Consensus:** A nonstationarity that is detected can be considered strong if it is detected by two or more detection methods of the same type (e.g., mean or variance/standard deviation or distribution). This represents consensus that a statistically significant nonstationarity in one metric occurs at a given point in a record. If consensus cannot be found for a given year or short period of time, then it is reasonable to discount it.

To analyze whether there is consensus among the methods, access the nonstationarity detection tab and view the middle graph on the right. This graph will display the individual nonstationarity tests and the corresponding nonstationarities detected in a given year.

2. **Robustness:** A statistically significant nonstationarity can be considered robust when tests targeting changes in two or more different statistical properties (mean, variance/standard deviation and/or overall distribution) are indicating a statistically significant nonstationarity (for example, the distributional LePagemo test, as well as the Lombard Mood and the CPM-Mood tests for changes in variance/standard deviation are identifying a nonstationarity). While a robust nonstationarity is not necessarily stronger, it represents a multifaceted change in the record. This can be taken into consideration when deciding which portion of the period of record to use to perform further analysis.

To determine whether the nonstationarity is robust, examine the tick marks associated with the year of the nonstationarity in the middle graph of the nonstationarity detection tab. If different colored tick marks are shown to indicate a nonstationarity in the same year, this indicates nonstationarities detected by tests that target changes in different statistical properties.

3. **Magnitude:** An identified nonstationarity is also associated with a given magnitude of change in the mean or standard deviation/variance in the datasets prior to and after the identified nonstationarity. Nonstationarities that are produced by greater changes in the statistical properties of the datasets before and after the identified nonstationarities may be important to take into consideration when performing subsequent analysis.

While presence of all three criteria is of value, consensus of multiple significant nonstationarities often connotes strong nonstationarity (Friedman, D. et al., 2018). For more information regarding nonstationarity detection and interpretation, please refer to Section 3.

1.7 Building Time Series Models

Time series models are powerful tools for characterizing data. They can be used in forecasting, error handling, interpretation, and decomposition. In addition, the analytic outputs of time series models can be immensely valuable in gaining deeper insight through model fitting. The three time series models used in the Time Series Toolbox are described as follows:

Time Series Model Choices

1. **Linear Models (TSLM):** Forecast the data values of interest, assuming that they have a linear relationship with time period of interest. This linear regression model also accounts for the trend and seasonality present in the data.



2. **Auto-Regressive Integrated Moving Average (ARIMA):** The most general class of models for forecasting a time series. It combines both recurring historical patterns (Autoregression) and overall trends (Moving Average) models as well as a differencing pre-processing step of the sequence to make the sequence stationary, called integration (I). The tool utilizes an Auto-ARIMA model that chooses the best ARIMA model according to either AIC, AICc, or BIC value. For more information on ARIMA, please read ARIMA Models, Chapter 8 (Hyndman, R. J., & Athanasopoulos, G., 2018).
3. **Exponential Smoothing (ETS):** A model that describes the data using weighted moving averages where the most recent data is weighted more than the historical data. These weights exponentially decrease as the observations get older. For more information on ETS models, please read Exponential Smoothing, Chapter 7 (Hyndman, R. J., & Athanasopoulos, G., 2018).

When choosing between ARIMA and ETS models, users should note if there is a trend and/or seasonality. If there are significant trends or seasonality, ETS may be a better option because it explicitly models these components. If there are little trends and few nonstationarities detected, ARIMA may be a better model as the prediction is a weighted sum of past observations (Hyndman, R. J., & Athanasopoulos, G., 2018).

Time Series Model Parameters

The three time series models offer a range of parameter options to customize the model. The tool automatically selects the default parameters for the TSLM and ETS models. For the ARIMA model, the tool uses optimal parameters based on the best fitting ARIMA model according to either AIC, AICc, or BIC value. However, users are also offered the option to manually edit model parameters. These parameters should be carefully reviewed before selecting. The following customizable parameters, obtained from Hyndman and Khandakar (2008), are defined below:

TSLM Parameters

- **Lambda:** Box-Cox transformation parameter. If lambda='auto', then a transformation is automatically selected using BoxCox.lambda. The transformation is ignored if NULL. Otherwise, data transformed before model is estimated.
- **Bias Adjustment:** If this parameter is checked (TRUE), the model uses an adjusted back-transformed mean for Box-Cox transformations. If transformed data is used to produce forecasts and fitted values, a regular back transformation will result in median forecasts. If biasadj is TRUE, an adjustment will be made to produce mean forecasts and fitted values.

ARIMA Model Parameters

- **Lambda:** Box-Cox transformation parameter. If lambda='auto', then a transformation is automatically selected using BoxCox.lambda. The transformation is ignored if NULL. Otherwise, data transformed before model is estimated.
- **Bias Adjustment:** If this parameter is checked (TRUE), the model uses an adjusted back-transformed mean for Box-Cox transformations. If transformed data is used to produce forecasts and fitted values, a regular back transformation will result in median forecasts. If biasadj is TRUE, an adjustment will be made to produce mean forecasts and fitted values.
- **Stationary:** If TRUE, restricts search to stationary models.
- **Seasonal:** If FALSE, restricts search to non-seasonal models.
- **Information Criterion:** Information criterion to be used in model selection.
- **Stepwise:** If TRUE, will do stepwise selection (faster). Otherwise, it searches over all models. Non-stepwise selection can be very slow, especially for seasonal models.
- **nmodels:** Maximum number of models considered in the stepwise search.



- **Approximation:** If TRUE, estimation is via conditional sums of squares and the information criteria used for model selection are approximated. The final model is still computed using maximum likelihood estimation. Approximation should be used for long time series or a high seasonal period to avoid excessive computation times.
- **Truncate:** An integer value indicating how many observations to use in model selection. The last truncate values of the series are used to select a model when truncate is not NULL and approximation=TRUE. All observations are used if either truncate = NULL or approximation = FALSE.
- **Seasonal Test:** This determines which method is used to select the number of seasonal differences. The default method is to use a measure of seasonal strength computed from an STL decomposition. Other possibilities involve seasonal unit root tests.
- **Allow Drift:** If TRUE, models with drift terms are considered.
- **Allow Mean:** If TRUE, models with a non-zero mean are considered.
- **Parallel:** If TRUE and stepwise = FALSE, then the specification search is done in parallel. This can give a significant speedup on multicore machines.
- **Number of Cores:** Allows the user to specify the number of parallel processes to be used if parallel = TRUE and stepwise = FALSE. If NULL, then the number of logical cores is automatically detected, and all available cores are used.

Autoregressive, Differencing, and Moving Average Components

- **d:** Order of first differencing. If missing, will choose a value based on test.
- **D:** Order of seasonal differencing. If missing, will choose a value based on the seasonal test method.
- **p:** Order of non-seasonal autoregressive
- **P:** Order of seasonal autoregressive
- **q:** Order of non-seasonal moving average
- **Q:** Order of seasonal moving average

ETS Model Parameters

- **Lambda:** Box-Cox transformation parameter. If lambda = 'auto', then a transformation is automatically selected using BoxCox.lambda. The transformation is ignored if NULL. Otherwise, data transformed before model is estimated.
- **Bias Adjustment:** If this parameter is checked (TRUE), the model uses an adjusted back-transformed mean for Box-Cox transformations. If transformed data is used to produce forecasts and fitted values, a regular back transformation will result in median forecasts. If biasadj is TRUE, an adjustment will be made to produce mean forecasts and fitted values.
- **Restrict:** If TRUE (default), the models with infinite variance will not be allowed.
- **Model:** A three-letter code indicating the model to be estimated using the ETS classification and notation. The possible inputs are “N” for none, “A” for additive, “M” for multiplicative, or “Z” for automatic selection. The default value of ZZZ ensures that all components are selected using the information criterion.
- **Damped:** If TRUE, use a damped trend (either additive or multiplicative). If NULL, both damped and non-damped trends will be tried and the best model (according to the information criterion ic) returned.
- **Additive Only:** If TRUE, will only consider additive models. Default is FALSE.
- **Optimization Criteria:** One of 'mse' (Mean Square Error), 'amse' (Average MSE over first nmse forecast horizons), 'sigma' (Standard deviation of residuals), 'mae' (Mean of absolute residuals), or 'lik' (Log-likelihood, the default).



- **Information Criteria:** Information criterion to be used in model selection.

Generalized Model Parameters

- **Confidence Interval:** A range of likely values for the estimates that will fall between a pair of values (upper and lower confidence values) around the mean. The user can select a 95%, 97%, or 99% confidence interval.
- **Forecast Horizon:** Length of time into the future for which forecasts are to be estimated.

Time Series Model Assessment

To assess the fit of the model chosen, there are two diagnostic plots provided to demonstrate residual behavior. These two models are: Residuals Over Time and Residuals Autocorrelation.

- **Residuals Over Time:** Shows the residuals or model errors over time. Positive values for the residual (on the y-axis) mean the prediction was too low, while negative values mean the prediction was too high; 0 indicates the guess was exactly correct. Ideally, this plot should look like random scatter around the line $y=0$ (Hyndman, R. J., & Athanasopoulos, G., 2018).
- **Residuals Autocorrelation:** Shows the autocorrelation or the relationship between the current observation and the previous ones. Ideally, the black bars eventually shrink within the blue bands.

In addition, a p-value from the Box-Ljung Test for the presence of statistically significant autocorrelation is provided. If this p-value is below .05, the null hypothesis of no autocorrelation is rejected (Laber & Shedden, 2017).

2 INTERACTING WITH THE TIME SERIES TOOLBOX

The Time Series Toolbox is designed to be both analytically robust and intuitive in use. It comprises multiple interfaces, each under a different subtab. This section of the user guide aims to direct the user through the application and address any potential questions the user may have on interaction with the tool. Future iterations of the guide will address any ad-hoc questions and concerns that are not addressed in Version 2.0.

The tool has six primary tabs on the top of the page: (1) Home, (2) Explore Data, (3) Model-Based Analysis, (4) Nonstationarity Detection, (5) Time Series Analysis, and (6) Help. Besides the Home, Time Series Analysis, and Help tabs—each tab contains subtabs that can be toggled by user selection.

The first tab of the tool (Home) provides text to introduce the user to the tool. The Explore Data page is designed to offer the user a place to upload and examine the dataset that will be used for analysis. The third tab, Model-Based Analysis, aims to provide a general analysis platform for the data—including trend and seasonality analysis. The fourth tab, Nonstationarity Detection, offers the user the ability to perform nonstationarity detection on the data. Furthermore, the Time Series Analysis tab allows the user to build time series models. The last tab, the Help tab, provides resources to the user, including model documentation and an email contact for any extraneous questions. The remainder of this section is dedicated to walking through each of these tabs within the tool.

2.1 Interactions with Graphs, Data Tables, and Time Filters

Graphs and tables appear throughout the tool to add additional context to the raw data and results. Since the raw information and analytical results are quantitative, data tables and charts are useful in presenting



and organizing the information. These items also allow the user to visualize the observations, which simplifies interpretation and drawing conclusions.

2.1.1 DATA TABLES

Several tabs of the Time Series Toolbox contain tables of the raw data or tables consisting of the data analysis findings. These graphs share an underlying foundation, and thus, users can interact with them in the same manner. Below are some of the main functionalities of the data table:

- **Copying the Data Table:** The upper left-hand corner displays an option to copy the data table; this table will then be copied into the user’s clipboard upon selection.
- **Exporting as CSV:** This button, also in the upper left-hand corner, allows the user to download the data as a CSV. The CSV document will appear in the user’s local download directory.
- **Exporting as a Microsoft Excel Document:** This button, also in the upper left-hand corner, allows the user to download the data as an Excel document. The Excel sheet will appear in the user’s local download directory.
- **Filtering and Sorting the Data:** The data can be sorted (highest to smallest data value, or by date value) by toggling the triangles to the right of the column headers. The data can also be filtered by date or data by controlling the sliders that appear as a user hovers over the “All” box (Figure 6).
- **Displaying X Entries:** A pop-down selection button is located on the top right for users to select the number of entries they would like to see on a given page.
- **Navigation:** Users can navigate between pages of data by selecting the page number or the previous/next buttons below the table.

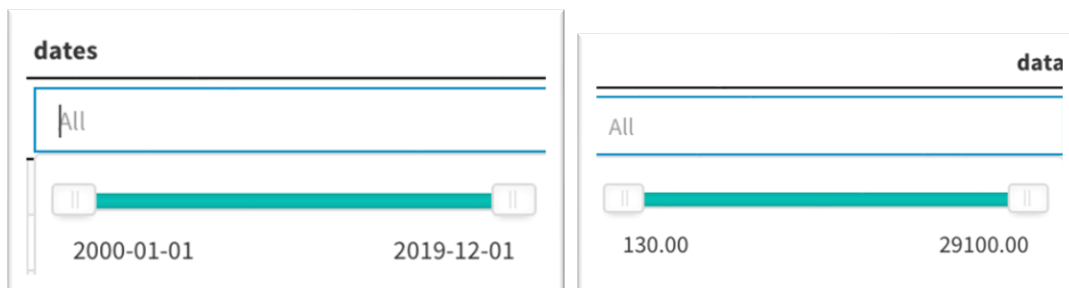


Figure 6: Data Table Sliders for Filtering Data



Raw Data Table

Copy CSV Excel Show 10 entries

dates	data
All	All
1901-01-01	26.6171
1901-02-01	26.5585
1901-03-01	26.8585
1901-04-01	27.3585
1901-05-01	27.7
1901-06-01	28
1901-07-01	27.8
1901-08-01	27.6
1901-09-01	27.8
1901-10-01	27.8

Showing 1 to 10 of 1,392 entries

Previous 1 2 3 4 5 ... 140 Next

Figure 7: Typical Data Table Shown Throughout the Application

Note: Some of the smaller tables do not have all these functionalities; they only have filtering and sorting features. These tables will not have the buttons on the top-left corner to export or download the graph.

2.1.2 GRAPHS

Several tabs of the Time Series Toolbox contain visualizations of the data, including Data Upload, Trend Analysis, Seasonality, Nonstationarity Detection, Breakpoint Analysis, and Time Series Analysis. These graphs serve to provide a visual image of the observations and results, which simplifies interpretation and drawing conclusions.

These visualizations use the same underlying technology—an R package called highcharter. Highcharter is an R wrapper for the Highcharts JavaScript library, a library used to create interactive web-based graphs. Since these graphics share the same underlying functionality, users can interact with these visualizations, or graphs, in the same manner.

Many of these actions are performed by selecting the icon (comprising three bars) on the top-right corner of the graphs (Figure 8). The different functionalities are also labeled in this figure. Below are some of the main export functionalities within the graphs:

- **Exporting Data Images from the Graph:** Users can export the data as an image in four different formats. Below are suggestions on which format may best suit the user's needs.



- **Download PNG Image (1):** Users can download the image in PNG image form. This is ideal for detailed graphics.
- **Download JPEG Image (2):** Users can download the image as a JPEG. This image format takes up very little storage space and is quick to upload or download.
- **Download PDF Document (3):** Users can download the image in a PDF document. The PDF will provide an electronic image of the graphic that looks the same as a printed document.
- **Download SVG Vector Image (4):** Users can download image as SVG vector image. The SVG image will remain crisp and clear at any resolution or size; however, file size will be large compared to the others listed.

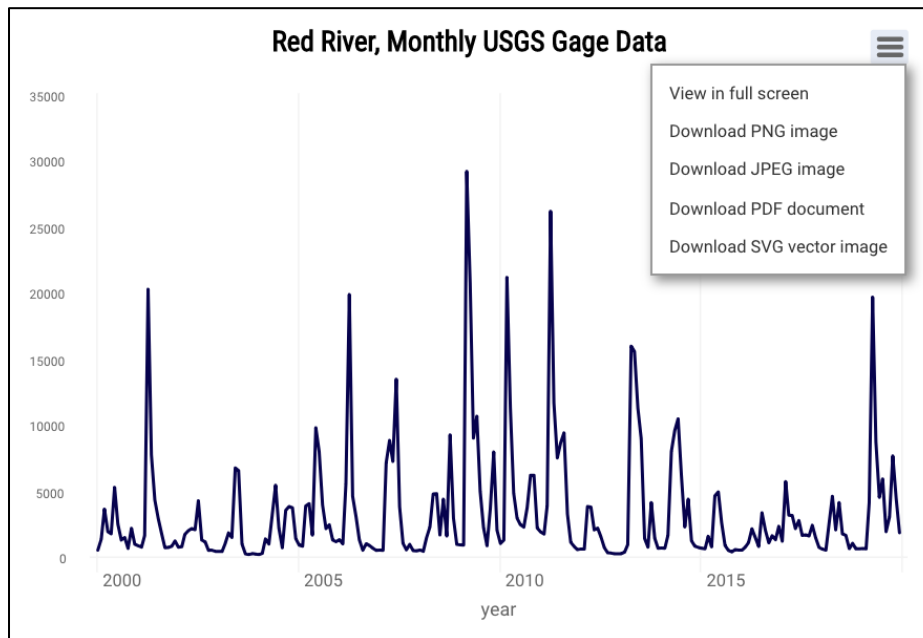


Figure 8: Visualization with Export Dropdown Options

If users have any issues with graph size (stretched or flattened graph), they can try expanding and changing browser size. Graphs will automatically adjust to these changes. In addition, users can zoom-in with a mouse pointer by dragging out a rectangle around the chart.

2.1.3 TIME SERIES FILTERING

Filtering by time can allow users to perform time series analysis and testing on segments of the data. This segmenting can be particularly useful for users after they run the change point and breakpoint analysis. From there, users can determine and highlight a segment of interest for future analysis. Time filters can be found in the Trend Analysis, Nonstationarity Detection, Breakpoint Analysis, and Time Series Analysis tabs. The dates can be adjusted in the text box or slider. Figure 9 provides an example of a time filter:

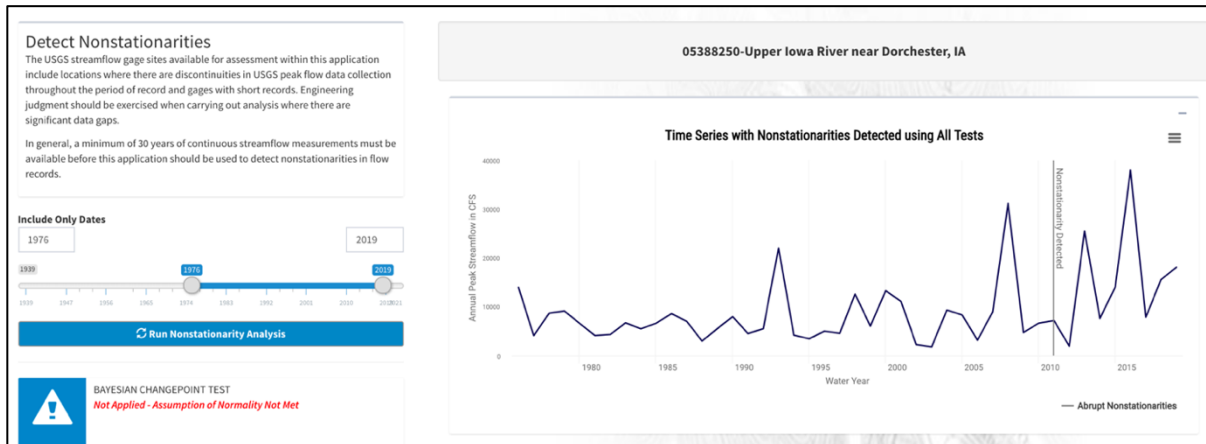


Figure 9: Timeframe Selection Featured in the Nonstationarity Detection Tab

2.2 Explore Data Interface

The Explore Data page allows the user to upload and examine uploaded data. The Explore Data panel consists of four subtabs: Data Upload, Data Summary, Magnificent Seven, and Seasonal Cycle. The user can navigate between the tabs by selecting the Explore Data tab on the top. The Data Upload subtab allows the user to select, upload, and visualize a dataset. The Data Summary, Magnificent Seven, and Seasonal Cycle subtabs help the user verify and preliminary analyze the data by highlighting summary statistics.

2.2.1 SUBTAB #1: DATA UPLOAD

This panel features two components—data upload and initial visualization of data. On the left-hand side of the application page is a panel comprising data source selection, text customization, and preprocessing options for user uploaded data. This panel is featured in Figure 1. For the user-uploaded data, the dataset must be in the format of a CSV file, with two data columns. The first column should consist of date values and the second of data values. Please refer to Section 1.3.1 for common errors for data formatting.

After the user selects the data sources, there is an option to customize the visualization of the data, including title, x-axis, and y-axis customization. These values will carry across throughout the visualizations used in tool. All of the customizations are optional; the default x-axis label is year.

Enter title (optional) - carries on to future tabs	
<input type="text" value="Red River, Monthly USGS Gage Data"/>	
Enter x-axis label (opt)	Enter y-axis label (opt)
<input type="text" value="year"/>	<input type="text"/>

Figure 10: Customization Options in the Upload Panel



Below these customization options are preprocessing options for user-uploaded data. The first preprocessing option determines how to handle missing values and carry them forward in the analysis. The second preprocessing option is data aggregation for data that has an observation frequency greater than one entry per year. If the dataset does not have more than one observation a year, this option will automatically be selected as “Do not aggregate data.” The third preprocessing option is detrending selection. These preprocessing sections are further described in Section 1.3.3 Data Preprocessing.

Figure 11: Preprocessing Options in the Data Upload

On the right side of the application, there is a time series visualization of the data. The visualization will appear on the top right after the user selects the “Upload Data” button (labeled “Load Data” for preloaded data). Once uploaded, two summary statistics will appear below the visualization (Missing Values and Total Data Points).




Figure 12: Summary Statistics Underneath the Time Series Plot

Within the visualization, the users can interact with the data by hovering over the line, which will reveal the specific data point (date, data). The visualization also allows users to export this graph as an image. Users can refer to Section 2.1.2 Graphs for more information regarding interacting with this visualization.

2.2.2 SUBTAB #2: DATA SUMMARY

Once the data is uploaded, the users can further explore the entries with the second subtab, the Data Summary subtab. This subtab includes a raw data table on the left. This data table allows the user to ensure the correct data was entered. Users can refer to Section 2.1.1 Data Tables for more information regarding interacting with data tables.

The right-hand side of the page displays more basic summary statistics, including start data, end date, frequency, missing values, and total entries.

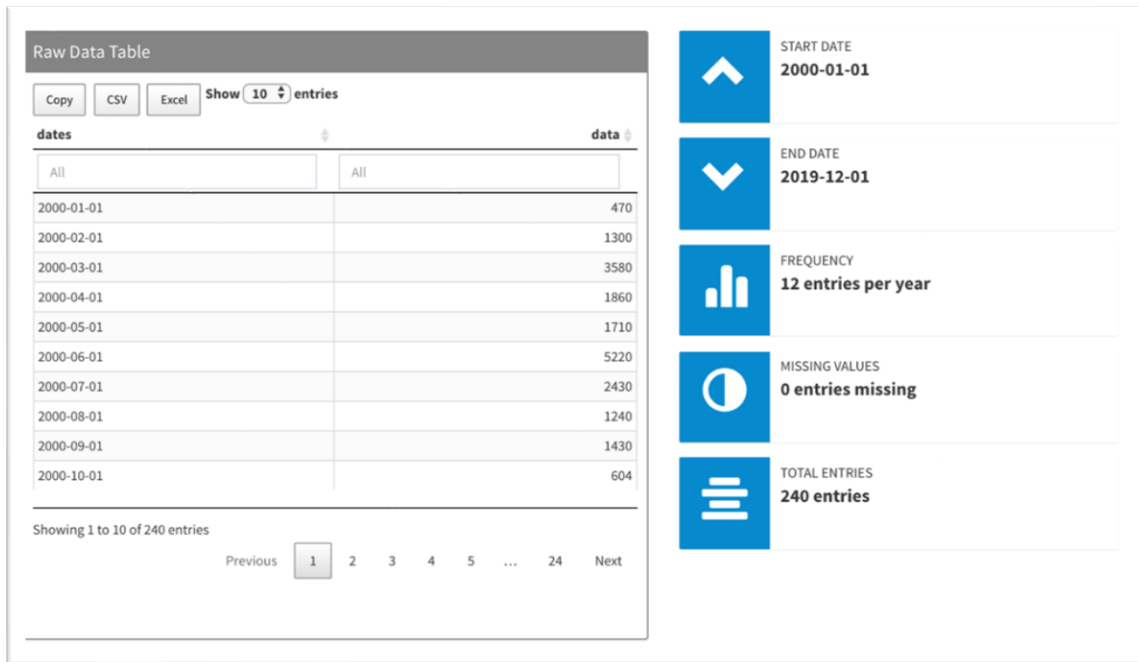


Figure 13: Data Summary Subtab

2.2.3 SUBTAB #3: MAGNIFICENT SEVEN ANALYSIS

This subtab displays the Magnificent Seven analysis for the dataset. This tab includes several information boxes with the metric (L-Mean, LCV (Coefficient of L-Variance), L-Skewness, L-Kurtosis, AR1, Amplitude and Phase) and the corresponding value. For more information about these statistics, please refer to Section 1.3.4 of the user guide—Data Exploration.

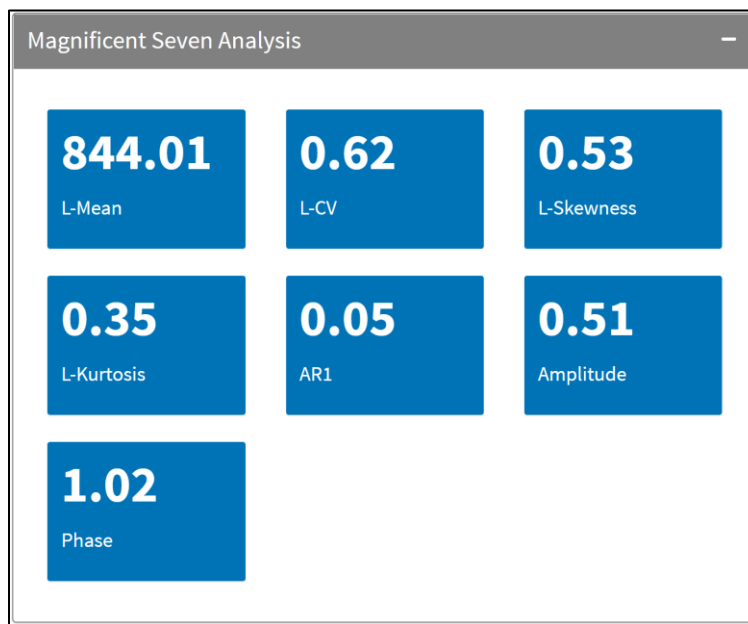


Figure 14: Magnificent Seven Information Boxes



2.2.4 SUBTAB #4: SEASONAL CYCLE

The seasonal cycle visualization and data table show variance in the data for a given month. Given monthly—or any smaller increment of—data the users can view the minimum, maximum, and average monthly values throughout all years. Please note, this tab is only available and shown if the data frequency is greater than one entry per year. Thus, this excludes annual data from being displayed in the Seasonal Cycle tab.




Figure 15: Seasonal Cycle Visualization

2.3 Model-Based Analysis

The Model-Based Analysis tab includes two model-based approaches to time series analysis: trend analysis and seasonality. These two models are represented as subtabs under the Model-Based Analysis main tab. Data uploaded in the Upload Data tab carries over to the Model-Based Analysis tab.

2.3.1 SUBTAB #1: TREND ANALYSIS

The Trend Analysis section detects and measures trends in time series data for generalized conclusions across the data series. On the left side of the section, there is a panel that displays information regarding the trend analysis, a time series filter, and a button to “Run Trend Analysis.” On the right-hand side, there is one main panel showcasing the visualization and trend hypothesis testing.

On the top of the right-hand tab, there is a visualization of the data with the two trend lines—traditional slope and sen’s slope—fit over the data. These trend values are also displayed below the graph on the left side table, the “Trend Line Coefficients” table. This table is originally open and can be collapsed with the “-” symbol on the right-hand side. When the table is expanded, users can sort the values in the table by selecting the arrows shown to the right of the column headers. To the right of the “Trend Line



Coefficients” table is the “Trend Hypothesis Test” table. This table displays the p-value from the t-Test, Mann-Kendall Test, and Spearman Rank-Order test. Beneath the table, there may be bullet points if there was a statistically significant trend detected by one of the tests. For more information regarding interacting with data tables, please see Section 2.1.1 Data Tables.



Figure 16: Trend Analysis Visualization

Furthermore, the trend analysis page will display infoboxes with text if there are any warnings associated with the data when running the trend analysis; if there are no warnings related to the current analysis, these infoboxes will not be displayed. These infoboxes are located to the left of the trend visualization and hypothesis tables. The first infobox emerges when there are less than 10 data entries in the record. The second warning message, the missing value infobox, will appear when there are any missing datapoints in the dataset.



DATA CHECK

WARNING: The period of record selected has less than 10 data points, the recommended number of observations. There are potential issues with the trends detected.

MISSING VALUES

WARNING: The period of record selected has missing data points. There are potential issues with the trends detected.

Figure 17: Data Checks for the Trend Analysis Tab

2.3.2 SUBTAB #2: SEASONALITY

The Seasonality subtab uses a series of statistical methods to identify and define seasonal patterns. The subtab will appear if the dataset qualifies for seasonality analysis; mainly, this requirement can be satisfied when the dataset is sub-annual.

The left panel describes seasonality and the methods used. It also allows users to define their decomposition analysis. This includes selecting the method for decomposition, which can either be the Moving Average Decomposition or Seasonal-Trend Decomposition by LOESS (STL). After running the seasonality decomposition, users can download this data on the bottom of the left panel. To the right of these options are the graphs of this decomposition, including the raw data, trend line data, seasonality, and random.

Define decomposition analysis:

Select method:

- Moving Average
- Seasonal-Trend Decomposition by LOESS (STL)

Figure 18: Decomposition Analysis Options




Figure 19: Trend Component of the Seasonality Decomposition




Figure 20: Random Component of the Seasonality Decomposition

Users can refer to Section 2.1.2 Graphs for more information regarding interacting with these visualizations. The visualizations are separate graphs; therefore, downloading, exporting as png, and other toolbox functionalities will be applied per graph.



The decomposition data can be downloaded by selecting “Download Seasonality Data.”

2.4 Nonstationarity Detection

The Nonstationarity Detection tab allows users to detect changing patterns in a dataset through two different subtabs—Nonstationarity Detection and Breakpoint Analysis. The Nonstationarity Detection subtab enables the user to apply a series of statistical tests to assess the stationarity of the inputted data. The second subtab, Breakpoint Analysis, identifies breakpoints in the data and displays them in a table.

2.4.1 SUBTAB #1: NSD SUBTAB

The first subtab is the Nonstationarity Detection subtab, which uses statistical testing to detect the presence of nonstationarities in the uploaded data. These tests examine the data for nonstationarities (or changes) in the data mean, variance, or distribution.

On the left, there is a panel that displays text information regarding stationary, the raw data frequency, timeframe filters, potential warning messages for limited or missing datasets, the method explorer, and options to configure parameters for the nonstationarity tests.

The grey button below the customization option allows for the configuration of the parameters for the nonstationarity tests. Once the button is selected, a new panel will pop up that allows users to configure the following parameters: CPM Methods Burn-In Period, CPM Methods Sensitivity, Bayesian Posterior Threshold, Energy Divisive Method Sensitivity, Pettitt Sensitivity, Bayesian Prior Likelihood, and Lombard Smooth Methods Sensitivity. The range for each parameter is listed on the left and right of the slider.

When the user is ready to run the analysis, they can select the “Run Nonstationarity Analysis” button, which will prompt three graphs to appear on the right panel. The first graph (shown in Figure 22) displays a visualization of the time series with any preprocessing from the Upload Data tab applied. Vertical lines show areas of nonstationarity.

Configure Testing Parameters

The below parameters are described in the Nonstationarity Detection Tool User Guide. Engineering judgment is required if parameters other than the default are used
For more information, please review the NSD Tool User Guide.

For these parameters, larger values will decrease the sensitivity of these tests, resulting in fewer nonstationarities detected.

CPM Methods Burn-In Period: 5 to 50 (Current: 20)

CPM Methods Sensitivity: 200 to 10,000 (Current: 1,000)

Bayesian Posterior Threshold: 0 to 1 (Current: 0.5)

Energy Divisive Method Sensitivity: 0.1 to 0.9 (Current: 0.5)

For these parameters, smaller values will decrease the sensitivity of these tests, resulting in fewer nonstationarities detected.

Pettitt Sensitivity: 0.01 to 0.15 (Current: 0.05)

Bayesian Prior Likelihood: 0.05 to 0.95 (Current: 0.5)

Lombard Smooth Methods Sensitivity: 0.01 to 0.15 (Current: 0.05)

Figure 21: Configuration of Parameters for Nonstationarity Tests




Figure 22: Nonstationarity Detection Graph (First NSD Graph)

The second graph displayed (shown in Figure 23) depicts the Statistical Tests Heatmap graphic. On the far left of the graph, it lists the 12 statistical tests being applied by the tool and displays a tick mark for each nonstationarity detected by the method applied. The color of the mark corresponds to the type of nonstationarity each test is targeted at detecting: mean-dark blue, variance/standard deviation-orange, or overall statistical distribution-light blue (discussed in Section 1.6.4). Nonstationarities detected by the Smooth Lombard tests are indicated by continuous yellow fill. As shown in **Error! Reference source not found.**, hovering over the tick marks displays the water year associated with the nonstationarity detected.

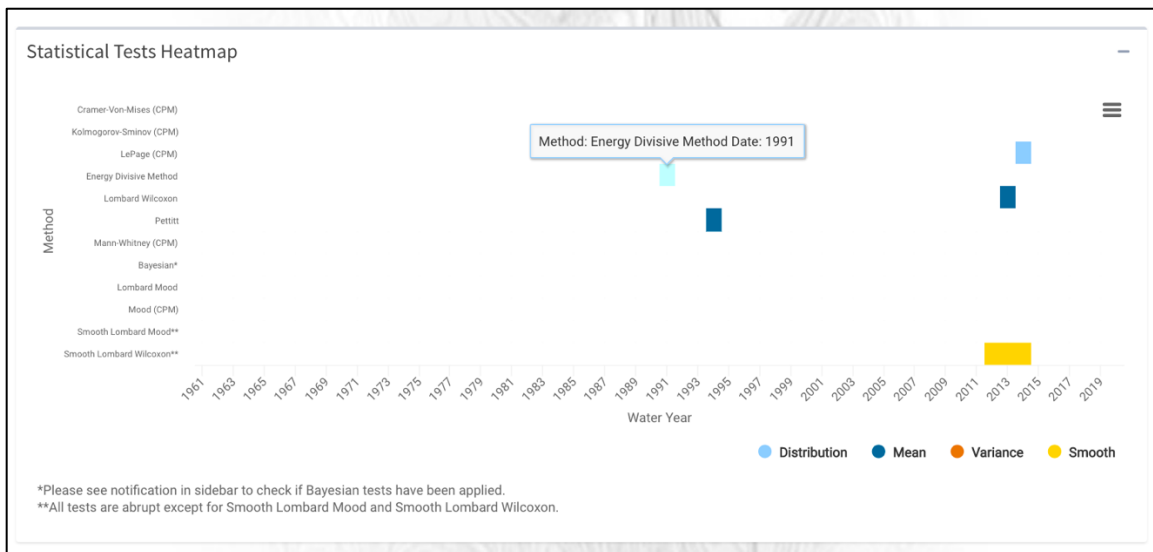


Figure 23: Nonstationarity Test Results Graph (Second NSD Graph)



The third graph displayed in Figure 24 depicts the segment mean, variance, and standard deviation for the subsets of data prior to and after each statistically significant nonstationarity detected. Each of these measures is presented with their associated units of measurement and is color coded to match the associated type of nonstationarity in the Statistical Tests Heatmap graphic. Note that, although the units for mean and standard deviation are the same, the units associated with sample variance are squared. This information can be used to determine the magnitude of changes in mean or variance/standard deviation associated with each of the nonstationarities detected. Hovering over the segments will show a hover-text of the date and the given segment value for that date.




Figure 24: Segment Statistics Graph (Third NSD Graph)

Users can refer to Section 2.1.2 Graphs for more information regarding interacting with these visualizations. These three graphs are all individual, so the application of toolbox buttons will apply to each of the graphs separately.

If the user selects preloaded gage data, additional metrics are available in the Nonstationary Detection tab. Primarily, an information box on the left panel (shown in Figure 25) will display the drainage information for that site. The user can also view site details and full data by selecting the “View Results and Gage Details” button (also shown in Figure 25). Within this pop-up information box, the Summary of Change Points tab details the identified nonstationarities by water year. By modifying how results are displayed using a series of checkboxes, the user can see summary tables indicating the statistical parameter which is changing for a given year and the individual tests which are showing evidence of nonstationarity. Section 3.3 contains guidance for how to interpret the results displayed. The user can also review the raw data being analyzed within the Full Data tab. Figure 26 displays the pop-up box that appears after the user selects that button. This box highlights the nonstationarities detected and shows the full original gage data.

Detect Nonstationarities

The USGS streamflow gage sites available for assessment within this application include locations where there are discontinuities in USGS peak flow data collection throughout the period of record and gages with short records. Engineering judgment should be exercised when carrying out analysis where there are significant data gaps.

In general, a minimum of 30 years of continuous streamflow measurements must be available before this application should be used to detect nonstationarities in flow records.

[View Results and Gage Details](#)

GAGE DRAINAGE AREA (SQM)

The gage has a drainage area of 84.8 square miles.

PLOT CHECK

If an axis does not line up, change the timeframe to start closer to the period of record

Figure 25: Gage Drainage Information Box and “View Results and Gage Details” Button on the Nonstationarity Detection Tab

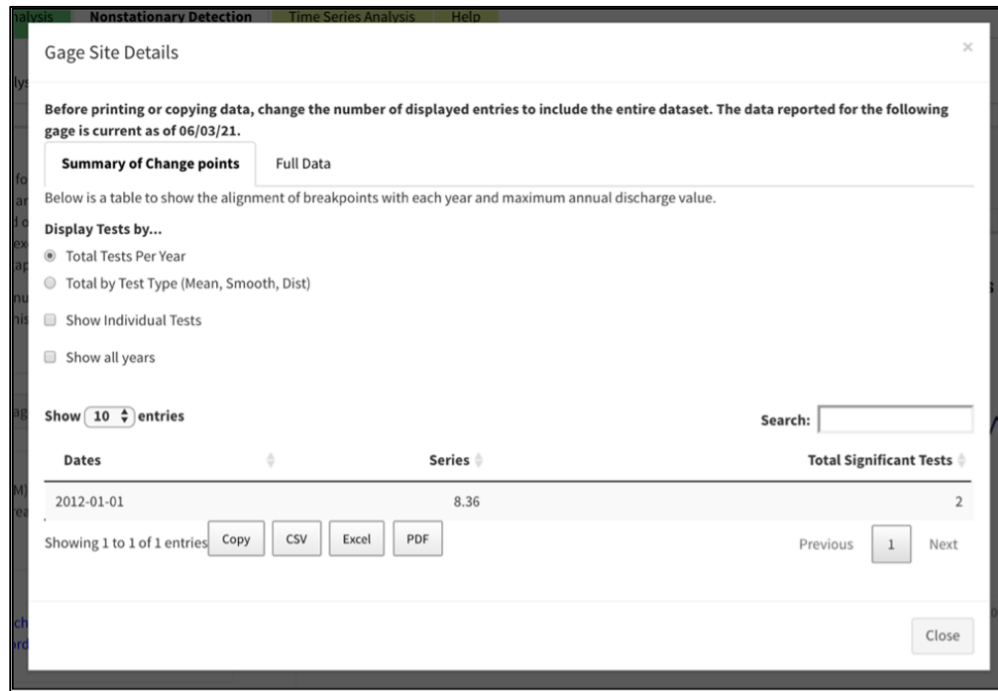


Figure 26: Additional Gage Details Pop-up on the Nonstationarity Detection Tab

Furthermore, the tool will display infoboxes with text if there are any warnings created in running the nonstationarity tests; if there are no warnings related to the current analysis, these infoboxes will not be displayed. These infoboxes can be seen in Figure 27. The first infobox is the missing value infobox that will appear when there are any missing datapoints in the dataset. The second infobox emerges when there are less than 30 entries in the record. This data length relates to the start and end date of the record, so missing values are included in the length. The third infobox refers to NSD testing parameters and will warn users if the default parameters have been changed—a selection made by the user. Finally, the Bayesian Change Point Test infobox, examines the normality of the data using the Anderson-Darling Test; the box will output a warning message if the assumption of normality is not met (i.e., the hypothesis of normality when the p -value is less than or equal to 0.05).

The user can also select and review the results derived from any one of the 12 nonstationarity detection methods (listed in Section 1.6.8) independently of the other statistical tests. Using the selection dropdown menu shown in Figure 28, the user can select and view the results for that individual test in the nonstationarity detection graph and the segment statistics graph. If non-default statistical test parameters (e.g., sensitivity parameters) are to be applied for a given test, the

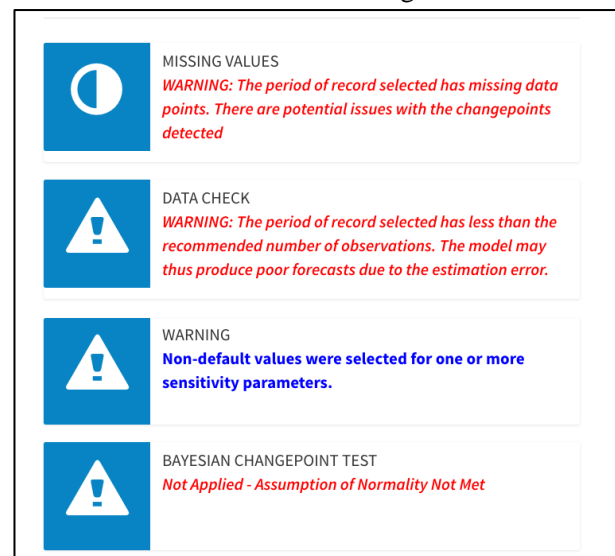


Figure 27: Possible Warning Infobox Messages That Appear in the Nonstationarity Detection Tab



Method Explorer should be used to examine the effects of adjustments to the parameters on the test results.

Method Explorer

If interested in individual test results, please select a test method to explore.

Note: Parameter values will be reset to default values after selecting new method.

All Tests ▼

****All tests are abrupt except for Smooth Lombard Mood and Smooth Lombard Wilcoxon**

Sensitivity Parameters

(Sensitivity parameters are described in the manual. Engineering judgment is required if non-default parameters are selected)

Configure NSD Test Parameters

Figure 28: Method Explorer Individual Test Selector

2.4.2 SUBTAB #2: BREAKPOINT ANALYSIS

The second subtab, Breakpoint Analysis, detects breakpoints in the data and displays the breakpoints as text in a raw data table and in a visualization format of a graph. As shown in **Error! Reference source not found.**, the section on the left shows a general text description of the breakpoint analysis, with a button below to detect the breakpoints. Below that button is the option to show regression lines between the optimized segments. Once the button is selected, a visualization and breakpoint table will be displayed on the right.

Users can view both the visualization and breakpoint table in the panel on the right. The visualization on the top of the right panel will display the plotted data and grey lines as the breakpoints identified in the graph. The users can also view the segment trend line by checking the “Show Segment Trend” checkbox (as seen in Figure 29**Error! Reference source not found.**). Users can refer to Section 2.1.2 Graphs for more information regarding interacting with this graph.



2.4.3

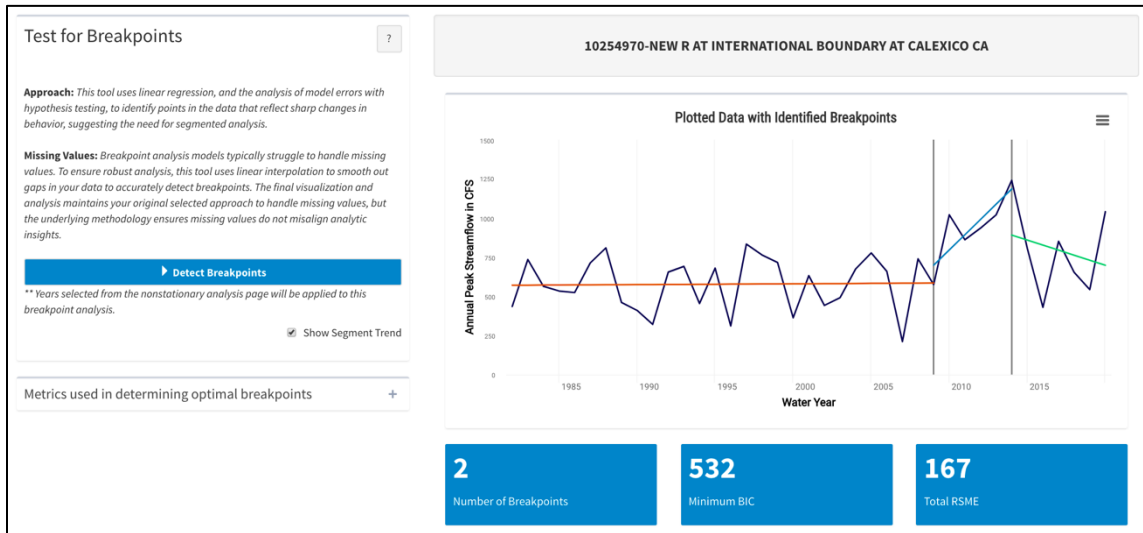


Figure 29: Breakpoint Visualization with Segment Linear Regression Lines Shown

Below this visualization is a data table that details segments in between each identified breakpoint. The following statistics are provided in the table: date range (start date to end date), trend coefficient, directionality, and intercept for the fitted regression line within that segment. This table can be copied to the clipboard with the “Copy” button.

Breakpoint Segment Details			
Date Range	Trend Coefficient	Intercept	Directionality
1982-01-01 to 2009-01-01	0.518	570.095	Positive
2009-01-01 to 2014-01-01	97.114	601.933	Positive
2014-01-01 to 2020-12-31	-32.25	923.714	Negative

Showing 1 to 3 of 3 entries Copy CSV Excel

Figure 30: Segment Details for Segments in Between Breakpoints

2.5 Time Series Modeling

For deeper inspection into time series data, the last section allows the user to build time series models. There are three model options: Time Series Linear Model, Auto Regressive Integrated Moving Average, and Exponential Smoothing. These options are displayed on the left-hand side of the page, in a column with text depicting the time series model and a button to run the time series analysis. Underneath the model selection, the user will also see options to customize time series model parameters, depending on what model is selected. More information surrounding these parameters can be found in Section 1.7.

The built models will be displayed on the right panel. This panel includes two subtabs—visualization and diagnostic plot subtabs—between which can be navigated on the top of the panel. The landing tab will be the fitted time series model with forecasted values in accordance with the selected forecast horizon. When



hovering over the data, the following characteristics will be displayed: the date, data value, and whether the data is the uploaded data, forecasted model, or the fitted model. When hovering over the forecasted data, red text will also be displayed. Users can refer to Section 2.1.2 Graphs for more information regarding interacting with the visualizations.

The second tab on the right panel is the diagnostic plots, which include Residuals Over Time and Residuals Autocorrelation plots. These graphs are not interactive like the other graphs in the tool.




Figure 31: Time Series Modeling Diagnostic Plots for Calexico, CA Streamflow Data



Furthermore, the Time Series Analysis tab will show infoboxes with text if there are any warnings associated with the data when running the time series models; if there are no warnings related to the current analysis, these infoboxes will not be displayed. These infoboxes are located to the left of the trend visualization and hypothesis tables. The first infobox is the missing value infobox that will appear when there are any missing datapoints in the dataset. The second infobox emerges when there are less than 30 data entries in the record.

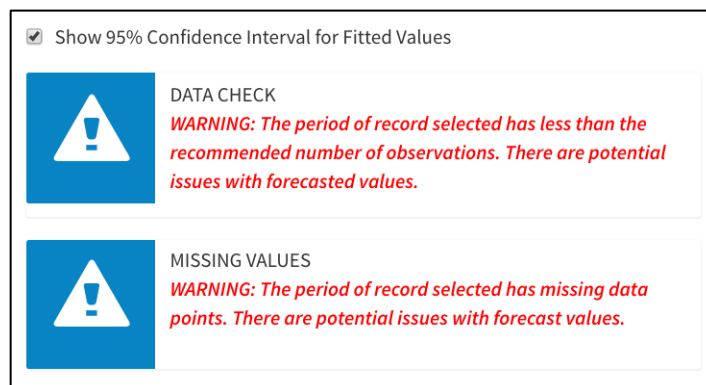


Figure 32: Possible Warning Infobox Messages That Appear in the Nonstationarity Detection Tab

3 EXAMPLE USE CASES

The following chapter discusses examples of how to conduct nonstationarity detection using the preloaded data in the Time Series Toolbox. This section explores an example where there is a priori knowledge that a nonstationarity may exist, one where there is no a priori knowledge, and another case where no nonstationarities are detected.

3.1 Simple Analysis (with a Priori Knowledge)

USGS ID 08089000 Brazos River Near Palo Pinto, Texas

This example demonstrates how to use the Time Series Toolbox to detect nonstationarities in an annual instantaneous peak streamflow record where there is a priori knowledge that a nonstationarity may exist in the record. The flow data used in this example was collected at the USGS gage site located on the Brazos River near Palo Pinto, Texas (USGS ID 08089000). Annual peak streamflow recorded on the Brazos River near Palo Pinto are known to be impacted by upstream regulation.

3.1.1 EXPLORATORY ANALYSIS

The first step in conducting time series analysis is to evaluate the data available and conduct an exploratory data analysis to better understand the data, in this case, the watershed where the gage is located. Figure 33 shows the annual instantaneous peak streamflow time series obtained from the USGS website. A visual examination of this time series suggests that there have been changes in the annual instantaneous peak streamflow record over the past 95 years. In particular, the values prior to the 1940s are on average higher than later years. Examination of the metadata associated with this record (available at: <http://wdr.water.usgs.gov/wy2013/pdfs/08089000.2013.pdf>) indicates that the construction of the Possum King Lake project was completed around 1941. This is also supported by the application of the USGS Peak Streamflow Qualification Code “6” (“Discharge affected by Regulation or Diversion”) to instantaneous peak streamflow data starting in 1941. Based on this information, a priori knowledge exists



that an abrupt change occurred in the early 1940s due to construction of the Morris Sheppard Dam, which impounds Possum King Lake. Therefore, the next step in the analysis is to formally test whether a nonstationarity exists in the annual instantaneous peak streamflow record observed at Palo Pinto in the early 1940s.

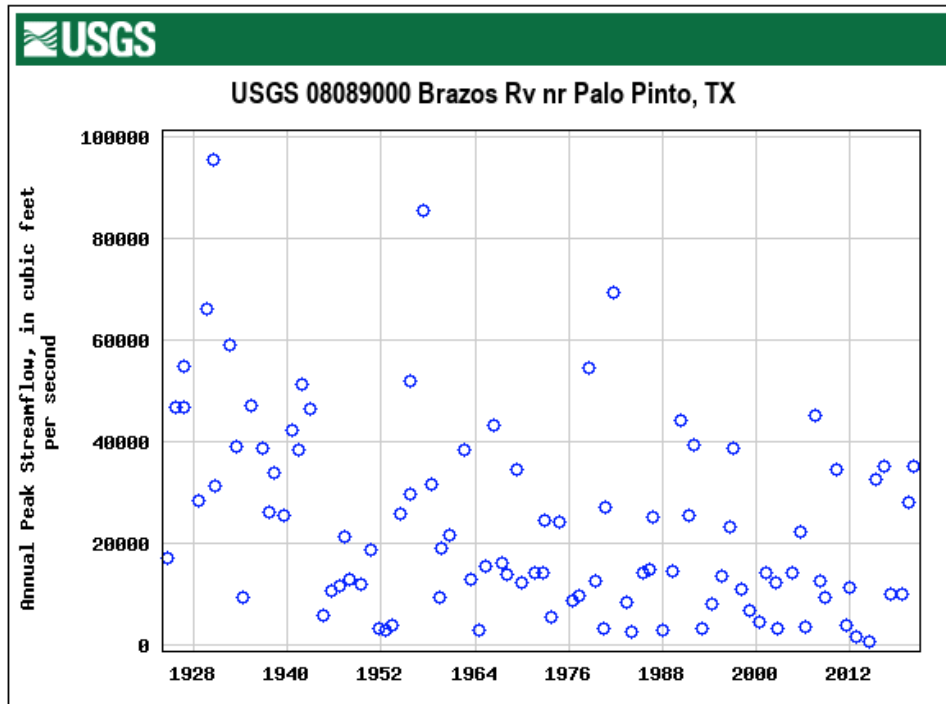


Figure 33: Annual Peak Streamflow Time Series for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000)

On the Explore Data tab, the user selects the data type to analyze, “Instantaneous Peak Stream,” site select search method, “Search by State,” and the state and site from the dropdown menus. The selected state is “Texas,” and the selected site is “08089000-Brazos RV NR Palo Pinto, TX.” After populating the dropdown menus, the user will select the “Load Data” button and then navigate to the Nonstationarity Detection page, which displays the Brazos River near Palo Pinto, TX nonstationarity detection analysis results (Figure 34). Because the annual instantaneous peak streamflow record at this gage site is continuous, the timeframe does not need to be adjusted before evaluating the results of the analysis.



Figure 34: Nonstationarity Detection Tab for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000)



3.1.2 NONSTATIONARITY DETECTION

As shown in Figure 34, the 12 nonstationarity detection methods contained in the tool collectively detected nonstationarities in two different years. To determine the types of changes detected by these methods, the user can click the View Results and Data button to view the statistical methods that detected the change (Figure 35).

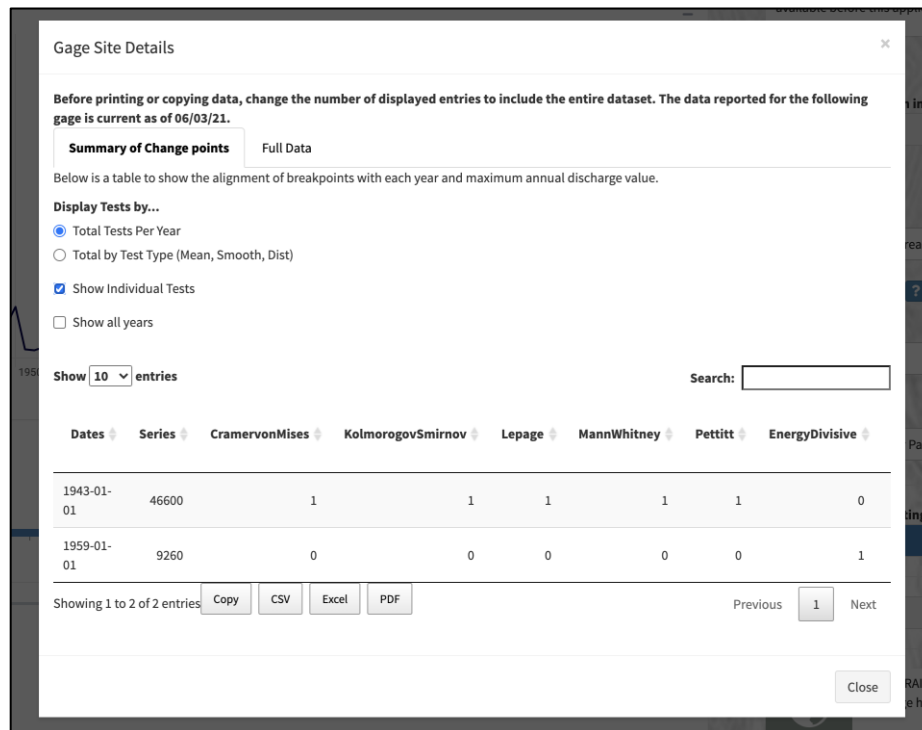


Figure 35: Gage Site Detailed Results for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000)

These results are also presented in the Statistical Tests Heatmap graphic. The color scheme associated with the tick marks in the Statistical Tests Heatmap graphic show that the detection methods are identifying nonstationarities driven by changes in the overall statistical distribution of the data and in the mean. A determination of the strength of the nonstationarities detected is carried out by evaluating whether they meet the criteria of consensus, robustness, and magnitude:

1. **Consensus:** The relative strength of a detected nonstationarity can be determined by looking at the level of consensus between different methods targeted at detecting the same type of nonstationarities (variance/standard deviation, mean, or overall distribution) in a data series. The nonstationarity detected in 1959 lacks consensus because it was only detected by a single method. The statistically significant nonstationarity detected in 1943 has consensus among three methods that detect a change in distribution and between two methods that detect a change in mean.
2. **Robustness:** The Statistical Tests Heatmap graphic in Figure 34 shows that multiple nonstationarity detection methods are detecting nonstationarities in the mean and the overall distribution of the dataset in the early 1940s. Consequently, the nonstationarity detected in the early 1940s can be considered representative of a “robust” change in the statistical properties of the flow dataset.



3. **Magnitude:** Using the Segment Statistics graphic displayed at the bottom of the Nonstationarity Detection tab (shown in Figure 36), the user can determine the magnitude of the shift in the mean of the annual instantaneous peak streamflow record associated with the 1943 strong, nonstationarity. The mean associated with annual peak flows collected between 1924-1942, as compared to 1944-2020 decreases from approximately 42,000 cfs to 19,000 cfs. A decrease in mean could be anticipated because the Possum King Lake flood control project was completed around 1941.

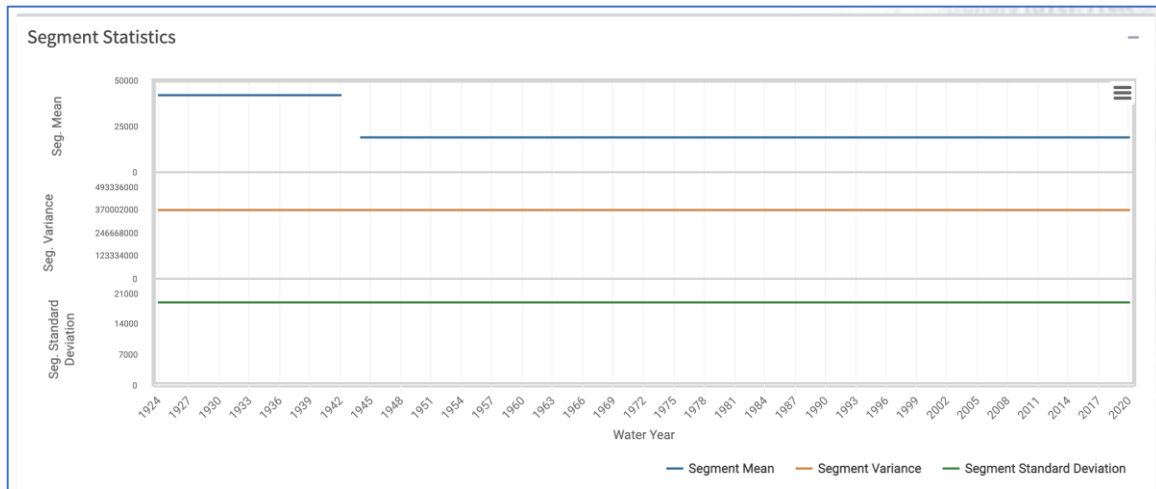


Figure 36: Segment Statistics Graphic for the Brazos River Near Palo Pinto in Texas (USGS ID 08089000)

In this example, the default sensitivity values facilitated the detection of the known nonstationarity in 1943. Consequently, the default sensitivity parameters were adopted for analysis. If the web tool had failed to detect the known nonstationarity in the period of record using the default sensitivity parameters, the user could take steps to incrementally adjust the sensitivity of the tests until the known nonstationarity is detected. Any other nonstationarities detected after these sensitivity parameters are adjusted are considered equally significant. If the user feels it is necessary to adjust the default parameters to identify a known nonstationarity based on prior knowledge, this would be documented, and the sensitivity parameters applied would be recorded.

Where there is a prior knowledge of a specific water management feature becoming operational, the user considers both local knowledge of when the homogenous operating period at the structure commenced and the nonstationarity that the tool is identifying when selecting a homogenous period of record for analysis.

3.1.3 MONOTONIC TREND ANALYSIS

The next portion of the analysis consists of assessing the period of record for monotonic trends. To assess monotonic trends within the record, the user can navigate to the Trend Analysis tab, where the tool performs multiple statistical tests to detect the presence of monotonic trends within the annual instantaneous peak streamflow data series.

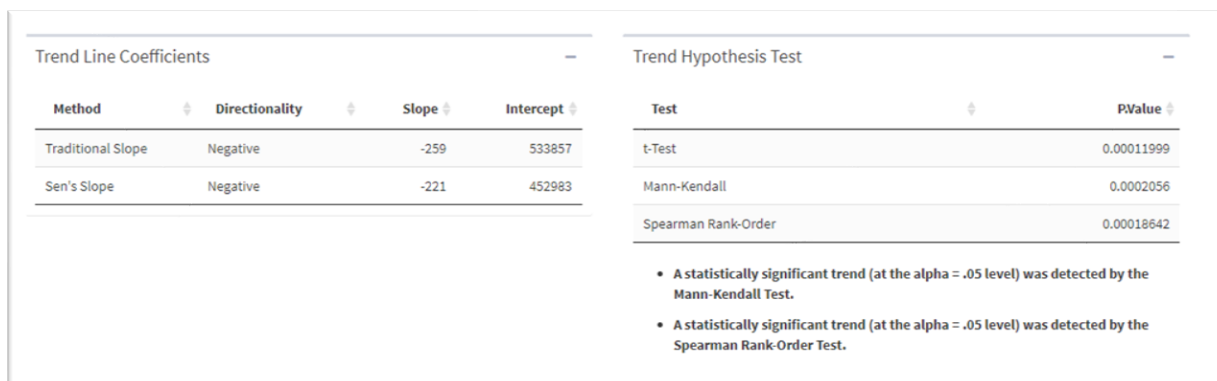


Figure 37: Monotonic Trend Analysis Results for the Full Period of Record 1924-2020

The results indicate that if the full period of record is analyzed there is a statistically significant decreasing trend in the data. This is consistent with the effect of the construction of a flood risk reduction project in 1941 (Morris Sheppard Dam).

3.1.4 APPLICATION TO HYDROLOGIC ANALYSIS

The results of the Nonstationarity Detection tab indicate that the statistical properties of the dataset collected prior to and post the construction of Morris Dam are different. Although the construction of the dam was completed in 1941, the Nonstationarity Detection tab identifies a nonstationarity in the dataset in 1943. This is likely because although the dam was constructed in 1941, it probably took a few years for the reservoir to fill and for “normal” operation to commence. For this reason, engineering judgment and a knowledge of reservoir operation/history is important when conducting hydrologic analysis.

3.2 Complex Analysis (Without a Priori Knowledge)

USGS ID 05082500 Red River of the North at Grand Forks, North Dakota

3.2.1 PROBLEM STATEMENT

Cases where nonstationarities are detected but cannot be directly attributed to a known source of non-homogeneity like the construction of a dam are more complex. An example of a case where nonstationarity is evident in the record without a readily attributable driver is the annual instantaneous peak streamflow for the Red River of the North at Grand Forks, North Dakota (USGS ID 05082500). Figure 38 shows the annual instantaneous peak streamflow time series obtained from the USGS. There is a substantial amount of drainage area between Grand Forks, North Dakota, and any upstream water control structures. Consequently, there is no prior knowledge indicating that there is a nonstationarity within the streamflow record at Grand Forks

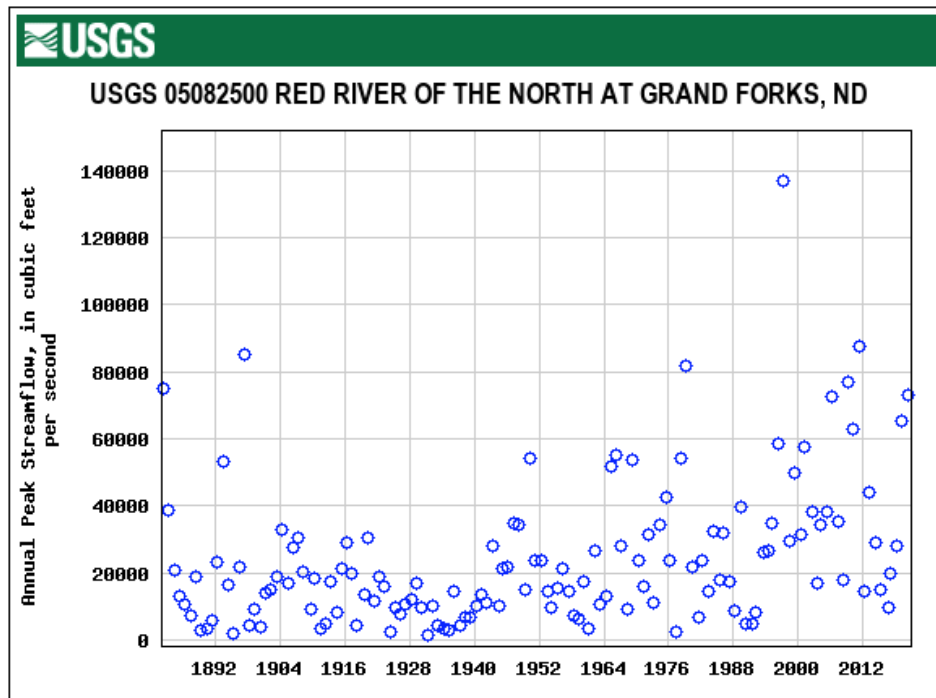


Figure 38: Annual Peak Streamflow Time Series for the Red River of the North at Grand Forks, ND (USGS ID 05082500)

The Time Series Toolbox was applied to check for nonstationarities within the streamflow record at Grand Forks, North Dakota.

3.2.2 SELECT A USGS SITE AND TIMEFRAME FOR NONSTATIONARITY DETECTION

On the Explore Data tab, the user selects the data type to analyze, “Instantaneous Peak Stream,” site select search method, “Search by State,” and the state and site from the dropdown menus. The selected state is “North Dakota,” and the selected site is “05082500- Red River of the North at Grand Forks, ND.” After populating the Explore Data tab, the user can navigate to the Nonstationarity Detection tab to analyze nonstationarities in the dataset. In this case, the timeframe does not need to be adjusted because the gage contains a continuous dataset for the entire period of record.



Figure 39: Nonstationarity Detection Tab for the Red River of the North at Grand Forks, ND (USGS ID 05082500)

3.2.3 NONSTATIONARITY DETECTION

As can be seen in the Time Series with Nonstationarities Detected graphic in Figure 39, the detection tab identified several nonstationarities, both abrupt (grey vertical lines) and smooth (orange shading), within the annual instantaneous peak streamflow record at Grand Forks.

As with the first example, hovering over the vertical bars in the Time Series with Nonstationarities Detected graphic indicates the exact year when an abrupt nonstationarity was detected. The Statistical Tests Heatmap graphic indicates which statistical tests are identifying nonstationarities in a given year. The tick marks in the graphic are color coded by the type of nonstationarity detected.

As seen in Figure 39, the Smooth Lombard Wilcoxon method also detects smooth changes in the mean of the annual instantaneous peak streamflow extending from 1936-2004 and the Smooth Lombard Mood method detects a smooth transition in sample variance/standard deviation of the annual instantaneous peak streamflow extending from 1993-1995. The output of the Smooth Lombard tests indicate that the statistics associated with the flow record at Grand Forks are in flux over these periods of time. Because the Smooth Lombard Mood result indicates a transitional period of five years or less, it can be compared directly to the abrupt nonstationarities indicated by the other techniques. It is analogous to an abrupt



nonstationarity in 1994. The longer period of nonstationarity (1936-2004) identified by the Smooth Lombard Wilcoxon alerts users that the stationarity assumption may not necessarily hold throughout this entire period.

It is notable that the Bayesian Change Point Test is not being applied because the dataset does not meet the required assumption of normality. There are three significant clusters of nonstationarity detection tests indicating changes in the statistical properties of the dataset: 1942, 1964-1965, and 1993-1995. The Energy Divisive Method also detects a statistically significant nonstationarity in 1912, but because this change point is only flagged by a single test it is not considered strong. An evaluation of the strength of the nonstationarities follows:

1. **Consensus:**

- a. **1942:** There is consensus between the Cramer-Von-Mises and the Mann-Whitney Test which indicates an abrupt change in mean circa 1942. In addition, the Smooth Lombard Wilcoxon indicates that the mean of the dataset is in flux starting in 1936.
- b. **1964-1965:** There is consensus between the LePage, Kolmogorov-Smirnov, and Energy Divisive methods which indicate changes in the distributional properties of flow datasets.
- c. **1994:** There is consensus between the Smooth Lombard Wilcoxon which indicates that the mean of the dataset is in flux between 1936 and 2004 and the Mann-Whitney Test which indicates an abrupt change in mean circa 1994.

2. **Robustness:** As can be seen in the Heatmap in Figure 39, the three nonstationarities can be considered robust.

- a. **1942:** The nonstationarity indicated in 1942 is flagged by tests targeted at detecting a change in mean (Mann-Whitney) and overall distribution (Cramer-Von-Mises).
- b. **1964-1965:** The nonstationarities detected between 1964-1965 are indicated by statistical tests that target changes in mean (Pettitt Test, Abrupt Wilcoxon, and Smooth Lombard Wilcoxon) and the overall statistical distribution of the datasets being assessed.
- c. **1994:** The nonstationarity detected in 1994 is indicated by statistical tests that target changes in mean, variance/standard deviation (Smooth Lombard Mood and Abrupt Lombard Mood), and overall distribution (Cramer-Von-Mises).

3. **Magnitude:** The magnitude of change in statistical properties can be gaged by using Segment Statistics graphic, located at the bottom of the Nonstationarity Detection tab (Figure 39).

- a. **1942:** The nonstationarity detected in 1942 amounts to an increase in mean from 16,000 to 20,000 cfs.
- b. **1964-1965:** The nonstationarities detected in the mid-1960s correspond to changes of about 7,000 cfs in the mean of the annual instantaneous peak streamflow.
- c. **1994:** The nonstationarity detected in 1994 corresponds to a change of 18,000 cfs in the mean and 13,000 cfs in the standard deviation of the annual instantaneous peak streamflow.

Because the nonstationarities detected in 1942, 1964-1965, and 1994 have consensus, can be considered robust, and represent appreciable changes in the mean and/or variance/standard deviation, the user can conclude that there are three, strong statistically significant nonstationarities within the dataset. Within the period 1964 to 1965 the majority of nonstationarities are detected in 1964. Consequently, 1964 can be adopted as the representative nonstationarity year for this period.



It is recommended that the default sensitivity values be adopted for analysis. Consequently, for this example application, default sensitivity parameters were applied. If the user chooses to alter the sensitivity parameters associated with the statistical tests executed by the web application, the new parameters could be recorded, and a description of how sensitivity parameters were modified could be included in report documentation.

3.2.4 MONOTONIC TREND ANALYSIS

The next portion of the analysis consists of analyzing the data for monotonic trends. To assess monotonic trends within this flow record, the user can navigate to the Trend Analysis tab, where the tool performs multiple statistical tests to detect the presence of monotonic trends in the annual instantaneous peak streamflow record, as seen in Figure 40. As you can see in the results below, the trend analysis shows a significantly significant increasing trend.

Table 2: Monotonic Trend Analysis: Red River of the North at Grand Forks, ND

(1)	(2)	(3)
Period of Record Analyzed	Statistically Significant Trend	Direction of Trend Detected
1882-2020	Yes	Increasing

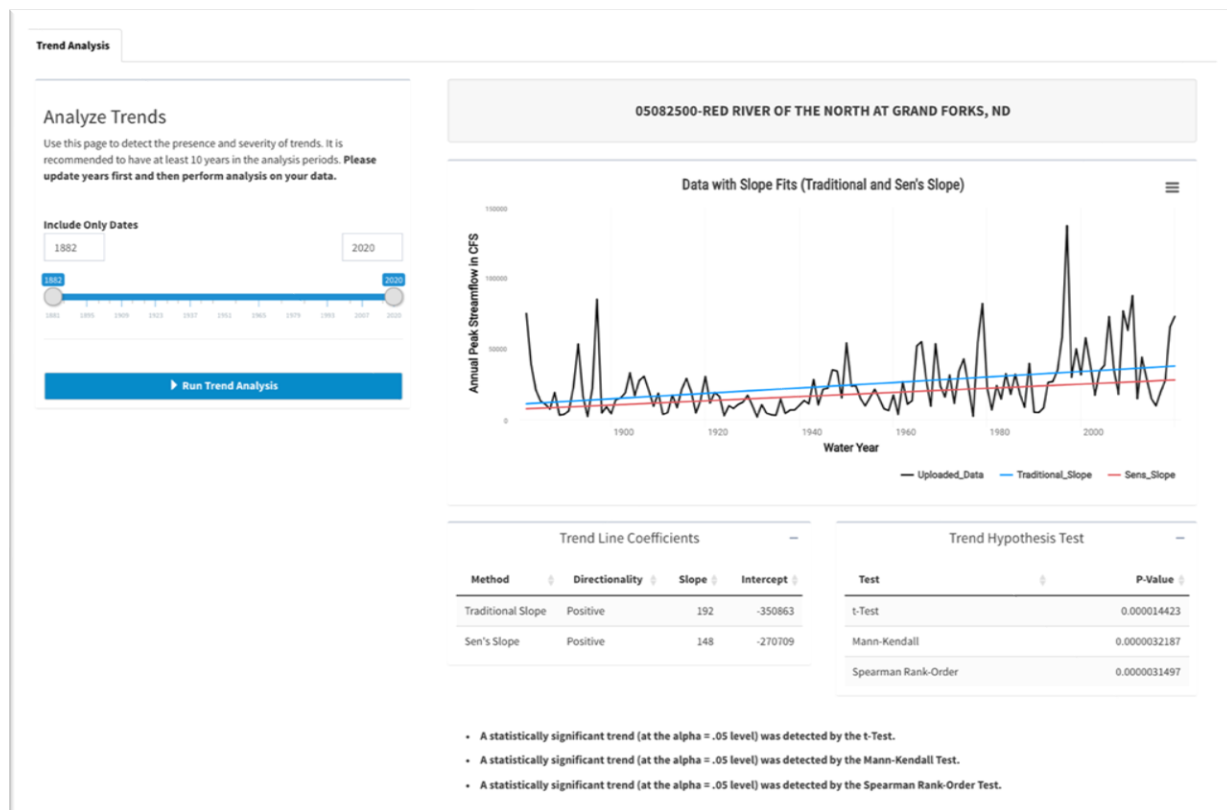


Figure 40: Trend Analysis Tab for the Red River of the North at Grand Forks, ND (USGS ID 05082500): 1882 -2020

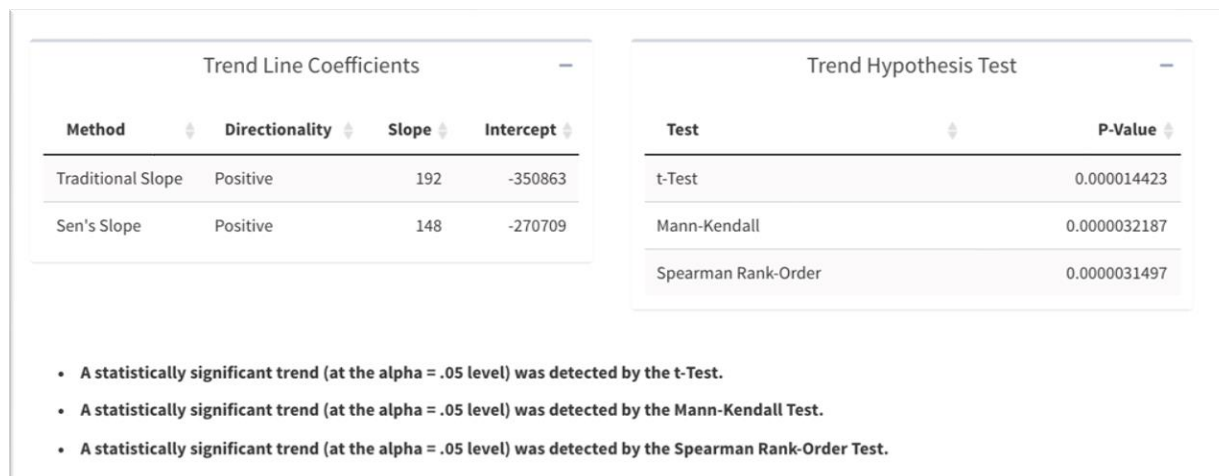


Figure 41: Trend Analysis Tab for the Red River of the North at Grand Forks, ND (USGS ID 05082500): Full Period of Record 1882-2020

3.2.5 APPLICATION TO HYDROLOGIC ANALYSIS

When a monotonic trend is detected in the dataset adopted for analysis, or when a long, transition period is detected by one of the smooth nonstationarity detection methods within the adopted period of record, the time series may still not meet the assumption of stationarity. If so, these findings are documented, and the user applies engineering judgment when carrying out hydrologic analysis in accordance with current USACE guidance: Engineer Manual (EM) 1110-2-1417, Flood-Runoff Analysis (1994), EM 1110-2-1413, and Hydrologic Analysis of Interior Areas (2018). Engineering judgment can consider the resiliency of the system when incorporating the range of results produced by conducting the sensitivity analysis described above in the adopted hydrologic design or study results.

Given the potential instability in this record, efforts to qualitatively evaluate what might be driving the nonstationarities detected in the flow record are advisable. The guidance presented in USACE Engineering and Construction Bulletin 2014-10, Guidance for Incorporating Climate Change Impacts to Inland Hydrology in Civil Works Studies, Designs, and Projects, could be applied to assess projected climate trends in the region. Temporal variation in land use practices could be studied using historic land use datasets and aerial imagery. Hydraulic structure inventories could be consulted to assess when both major and minor hydraulic structures were installed throughout the study area.

3.3 No Statistically Significant Nonstationarity Detected

USGS ID 05388250: Upper Iowa River Near Dorchester, Iowa

3.3.1 PROBLEM STATEMENT

This example assesses the stationarity assumption as it applies to the annual instantaneous peak streamflow dataset recorded along the Upper Iowa River near Dorchester, Iowa (USGS ID 05388250). Figure 42 shows the annual instantaneous peak streamflow time series obtained from the USGS for gage 05388250. There is no prior knowledge indicating that there might be a nonstationarity within the streamflow record at Dorchester, Iowa.




Figure 42: Annual Peak Streamflow Time Series for the Upper Iowa River Near Dorchester, Iowa (USGS ID 05388250)

3.3.2 SELECT A USGS SITE AND TIMEFRAME FOR NONSTATIONARITY DETECTION

On the Data Upload subtab of Explore Data, the user selects the data type to analyze, “Instantaneous Peak Stream,” site select search method, “Search by State,” and the state and site from dropdown menus 1-5. The selected state is “Iowa,” and the selected site is “Upper Iowa River Near Dorchester, IA.” After populating the dropdown menus and selecting the upload data button, the user can navigate to the Nonstationarity Detection tab, to find the Upper Iowa River near Dorchester, IA nonstationarity detection analysis results (Figure 43). In Figure 43, the annual instantaneous peak streamflow time series in the Time Series with Nonstationarities Detected graphic shows missing data between 1941 and 1976.




Figure 43: Nonstationarity Detection Tab for the Upper Iowa River Near Dorchester, Iowa (USGS ID 05388250)—Without Timeframe Adjustment

Using the timeframe filter, modify the timeframe selected for analysis to the period with continuous flow data: 1976-2020. Note that changing the timeframe will impact the nonstationarities detected by the tool.



3.3.3 NONSTATIONARITY DETECTION

Now that the timeframe and site have been selected, the results of the detection analysis can be assessed. As shown in Figure 44, only a single nonstationarity was detected by the tool in 2011 by one method: Lombard Wilcoxon. Because only one method is indicating a change point this implies a lack of consensus and robustness, thus the detected nonstationarity cannot be considered strong.




Figure 44: Nonstationarity Detection Tab for the Upper Iowa River Near Dorchester, Iowa (USGS ID 05388250)—with Timeframe Adjustment



3.3.4 MONOTONIC TREND ANALYSIS

As can be seen in Figure 45, there is no statistical evidence of a monotonic trend in the annual instantaneous peak streamflow recorded on the Iowa River near Dorchester, IA.

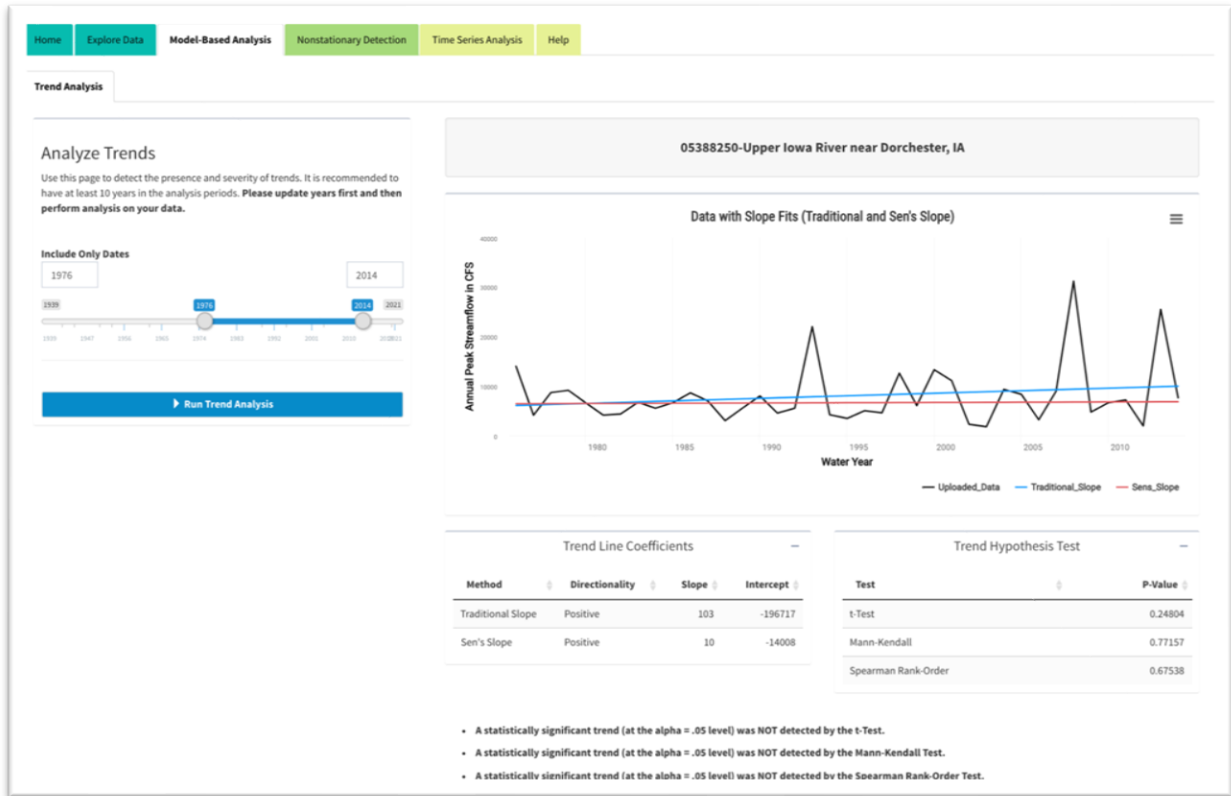


Figure 45: Trend Analysis Tab for the Upper Iowa River Near Dorchester, Iowa (USGS ID 05388250)

Because there is no evidence of nonstationarity in the streamflow record observed near Dorchester, Iowa inland hydrologic analysis will be evaluated in accordance to current USACE guidance: Engineer Manual (EM) 1110-2-1417, Flood-Runoff Analysis (1994), EM 1110-2-1413, and Hydrologic Analysis of Interior Areas (2018).



ADDITIONAL REFERENCES

- Adhikari, Ratnadip & K. Agrawal, R. (2013). An Introductory Study on Time Series Modeling and Forecasting. 10.13140/2.1.2771.8084.
- Apurv, T., & Cai, X. (2019). Evaluation of the Stationarity Assumption for Meteorological Drought Risk Estimation at the Multidecadal Scale in Contiguous United States. *Water Resources Research*, 55. <https://doi.org/10.1029/2018WR024047>
- Archfield A., S & Kennen, Jonathan & Carlisle, Daren & Wolock, David. (2013). An Objective and Parsimonious Approach for Classifying Natural Flow Regimes at a Continental Scale. *River Research and Applications*. 30. 10.1002/rra.2710.
- Bindoff, N. L., J. Willebrand, V. Artale, A. Cazenave, J. Gregory, S. Gulev, K. Hanawa, C. Le Quéré, S. Levitus, Y. Nojiri, C. K. Shum, L. D. Talley, and A. Unnikrishnan. (2007). “Chapter 5, Observations: Oceanic Climate Change and Sea Level.” In: *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change* (S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. B. Averyt, M. Tignor, and H. L. Miller, eds.). Cambridge, United Kingdom, and New York, NY: Cambridge University Press.
- Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. J. (1990). STL: A Seasonal-Trend Decomposition Procedure Based on Loess. *Journal of Official Statistics*, 6(1), 3–33.
- Climate Change Science Program. (2009). *Synthesis and Assessment Product 4.1: Coastal Sensitivity to Sea Level Rise: A Focus on the Mid-Atlantic Region. A Report by the U.S. Climate Change Program and the Subcommittee on Global Change Research*. [J. G. Titus (Coordinating Lead Author), E. K. Anderson, D. Cahoon, S. K. Gill, R. E. Thieler, J. S. Williams (Lead Authors)]. Washington, DC: U.S. Environmental Protection Agency. <http://www.climatechange.gov/Library/sap/sap4-1/final-report/default.htm>
- Deputy Secretary of Defense. (2003). *Ensuring Quality of Information Disseminated to the Public by the Department of Defense*, Department of Defense: Washington, DC. http://www.usace.army.mil/Portals/2/docs/authorization_form.pdf
- Erdman, C., & Emerson, J. W. (2007). bcp: An R Package for Performing a Bayesian Analysis of Change Point Problems. *J. Stat. Software*, 1–13.
- Friedman, D., J. Schechter, Sant-Miller, A. M., C. Mueller, G. Villarini, K. D. White, and B. Baker. (2018), *U.S. Army Corps of Engineers Nonstationarity Detection Tool User Guide*. U.S. Army Corps of Engineers (USACE): Washington, DC.
- Grigorios Papageorgiou, Stuart W. Grant, Johanna J. M. Takkenberg, and Mostafa M. Mokhles. Statistical Primer: How to Deal with Missing Data in Scientific Research. *Interactive CardioVascular and Thoracic Surgery*, Volume 27, Issue 2. August 2018. Pages 153–158. <https://doi.org/10.1093/icvts/ivy102>
- Hosking, J. R. M. (1990). L-Moments: Analysis and Estimation of Distributions Using Linear Combinations of Order Statistics. *Journal of the Royal Statistical Society, Series B*, 52, 105-124.
- Hosking, J. R. M. (1996). Fortran Routines for Use with the Method of L-Moments, Version 3. Research Report RC20525, IBM Research Division, Yorktown Heights, NY.



Hyndman, R. J. and Khandakar, Y. (2008). Automatic Time Series Forecasting: The Forecast Package for R. *Journal of Statistical Software*, 26(3).

Hyndman R., Athanasopoulos G., Bergmeir C., Caceres G., Chhay L., O'Hara-Wild M., Petropoulos F., Razbash S., Wang E., Yasmeeen F. (2020). *Forecast: Forecasting Functions for Time Series and Linear Models*. R Package Version 8.13. <https://pkg.robjhyndman.com/forecast/>

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*. Heathmont, Vic.: OTexts.

Intergovernmental Oceanographic Commission. (1985). *Manual on Sea Level Measurement and Interpretation, Volume I*. Intergovernmental Oceanographic Commission Manuals and Guides 14. <http://unesdoc.unesco.org/images/0006/000650/065061eb.pdf>

Intergovernmental Panel on Climate Change (IPCC). (2007a). *Climate Change 2007: The Physical Science Basis, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change* (S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. B. Averyt, M. Tignor, and H. L. Miller (eds.)). Cambridge, United Kingdom, and New York, NY: Cambridge University Press. <http://ipcc-wg1.ucar.edu/wg1/wg1-report.html>

Intergovernmental Panel on Climate Change (IPCC). (2013). *Climate Change 2013: The Physical Science Basis, Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)). Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 1535 pp, doi:10.1017/CBO9781107415324.

James, N. A., & Matteson, D. S. (2015). ecp: An R Package for Nonparametric Multiple Change Point Analysis of Multivariate Data. *Journal of Statistical Software*, 62.7.

Kumar, P., Lahiri, B., MM, N., & Sangma, S. (2022). Trend analysis and Changepoint Detection of Monthly, Seasonal and Annual Climatic Parameters in The Garo Hills of Northeast India. *Research Square*, <https://doi.org/10.21203/rs.3.rs-1412182/v1>

Laber, E. B., & Shedden, K. (2017). Statistical Significance and the Dichotomization of Evidence: The Relevance of the ASA Statement on Statistical Significance and p-Values for Statisticians. *Journal of the American Statistical Association*, 112(519), 902-904. doi:10.1080/01621459.2017.1311265.

Lund, R., X. L. Wang, Q. Q. Lu, J. Reeves, C. Gallagher, and Y. Feng, 2007: Changepoint Detection in Periodic and Autocorrelated Time Series. *J. Climate*, 20, 5178–5190, <https://doi.org/10.1175/JCLI4291.1>.

Militino, Ana F., Mehdi Moradi, and M. Dolores Ugarte. (2020). On the Performances of Trend and Change-Point Detection Methods for Remote Sensing Data. *Remote Sensing* 12(6): 1008. <https://doi.org/10.3390/rs12061008>

Milly, P., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., & Stouffer, R. J. (2008). Stationarity Is Dead: Whither Water Management? *Science*, 319(5863), 573-574. doi:10.1126

Northeast Power Coordinating Council. (2013b). *NPCC2 Climate Risk Information 2013: Climate Methods Memorandum*, New York City Mayor's Office of Sustainability: New York, NY.



http://www.nyc.gov/html/planyc2030/downloads/pdf/NPCC2_Climate%20Methods%20Memo%20random_2013.pdf

Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P. M. Midgley (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 1535 pp. <https://www.ipcc.ch/report/ar5/wg1/>

Pohlert, T. (2015). Trend: Non-Parametric Trend Tests and Change-Point Detection. R Package Version 0.0.1. doi:10.13140.

Quessy, J.-F., Favre, A.-C., Said, M., & Champagne, M. (2011). Statistical Inference in Lombard's Smooth-Change Model. *Envirometrics*, 882-93.

R Core Team. (n.d.). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing.

Ross, G. J. (2015). Parametric and Nonparametric Sequential Change Detection in R: The Cpm Package. *Journal of Statistical Software*, 66.3.

Ryberg, K. R., & Vecchia, A. V. (2012). *Waterdata—An R Package for Retrieval, Analysis, and Anomaly Calculation of Daily Hydrologic Time Series Data, Version 1.0*. U.S. Geological Survey. Bismarck, N.D.: U.S. Geological Survey. Retrieved September 18, 2015. <http://pubs.usgs.gov/of/2012/1168/of12-1168.pdf>

Tan, X. and T.Y. Gan (2015). Nonstationary Analysis of Annual Maximum Streamflow of Canada. *J. Climate*, 28, 1788–1805. <https://doi.org/10.1175/JCLI-D-14-00538.1>

USACE. (2018). *Hydrologic Analysis of Interior Areas*. Washington, DC: USACE. Retrieved June 2, 2022. https://www.publications.usace.army.mil/Portals/76/Users/227/19/2019/EM_1110-2-1413.pdf?ver=2018-09-13-125346-687

USACE. (1989). *Engineer Circular (EC) 1105–2–186 (Expired): Guidance on the Incorporation of Sea Level Rise Possibilities in Feasibility Studies*. Washington, DC: USACE.

USACE. (1994). *Flood-Runoff Analysis*. Washington, DC: USACE . Retrieved September 18, 2015. http://www.publications.usace.army.mil/Portals/76/Publications/EngineerManuals/EM_1110-2-1417.pdf

USACE. (2011). *Climate Change Adaptation Policy Statement*. June 3, 2011. <http://www.corpsclimate.us/docs/USACEAdaptationPolicy3June2011.pdf>

USACE. (2013a). *Engineer Regulation 1100-2-8162, Incorporating Sea-Level Change in Civil Works Programs*. USACE: Washington, DC. http://www.publications.usace.army.mil/Portals/76/Publications/EngineerRegulations/ER_1100-2-8162.pdf

USACE. (2014). *Engineer Technical Letter 1100-2-1, Procedures to Evaluate Sea Level Change: Impacts, Responses, and Adaptation*. USACE: Washington, DC. http://www.publications.usace.army.mil/Portals/76/Publications/EngineerTechnicalLetters/ETL_1100-2-1.pdf



USACE. (Forthcoming). *Guidance for Detection of Nonstationarities in Annual Maximum Discharges*. Washington, DC: USACE.

USACE. (2016). *Modified Procedure for Computing Relative Sea Level Change in Areas Exhibiting Rapid Relative Sea Level Change*. Washington, DC: USACE.

Velicer, Wayne & Fava, Joseph. (2003). Time Series Analysis. 10.1002/0471264385.wei0223.

Yaman, B., & Ertuğrul, M. (2020). Change-point detection and trend analysis in monthly, seasonal and annual air temperature and precipitation series in Bartın province in the western Black Sea region of Turkey. *Geology, Geophysics and Environment*, 46(3), 223. <https://doi.org/10.7494/geol.2020.46.3.223>

Zervas, C. (2009). *Sea Level Variations of the United States 1854–2006*. NOS CO-OPS 053. Silver Spring, MD: Center for Operational Oceanographic Products and Services, National Ocean Service, National Oceanic and Atmospheric Administration.

Zhou, Y., Ruan, G., Xu, C.-Y., Xiong, L., Jain, S. K., & Li, L. (2022). Detection and attribution of Norwegian annual precipitation variability related to teleconnections. *Earth and Space Science*, 9, e2021EA001857. <https://doi.org/10.1029/2021EA001857>



APPENDIX 1: GLOSSARY OF STATISTICS TERMS

Bayesian: In contrast to “Frequentist” statistics, Bayesian statistics interprets probability as a quantity assigned to a state of belief. In the Bayesian view, a probability can be assigned to a hypothesis, which is known as the prior.

Monotonic Trend: A variable increases or decreases consistently over time. This does not imply a linear trend.

Parametric: From a branch of statistics that assumes the sample data comes from a population that follows a probability distribution based on a fixed set of parameters.

Posterior Probability: This is unique to Bayesian statistics. It is the outcome of a Bayesian model. A Bayesian assigns a prior probability to a hypothesis, then uses the data to inform the model, before outputting the new state of belief around an event.

Stationarity: The assumption that the statistical characteristics of hydrologic time series data are constant through time.

Type I Error: The probability of accepting an alternative hypothesis when the null hypothesis is true.



APPENDIX 2: VERSION CHANGE LOG

Version 1.0.4 Change Log

- Data in the tool was updated to reflect USGS information as of 06/03/2021.
- Added the ability to view raw data in tabular format and export for use outside of tool
- Added the ability to view model results in tabular format and export for use outside of tool
- Updated user interface to maximize room for data visualizations
- Added the ability to change the Bayesian Prior Likelihood parameter in the tool for calculating Bayesian Change points
- Added a Site Selector tab for users to view physical information about the site
- Added summary infoboxes to Trend Analysis page
- Added dynamic titles for the Method Explorer page
- Changed date formats to water year only formats for visualizations on the Detection and Method Explorer page
- Added the full test names to the heat map on the Detection page
- Added separation borders for the Segment Statistics graph
- Updated the y-axis on the Segment Statistics graph so that all values will display on the left side of the visualization
- Updated colors on the Detection, Trend Analysis, and Method Explorer graphs to be Section 508 C compliant
- Added notification to alert users when BCP tests were/were not applied
- Added help text to remind users of abrupt vs smooth tests on heatmap
- Added water summary and gage data links to Site Selector tab
- Added new column to table in the Trend Analysis tab to note the directionality of the trends

Version 2.0 Change Log

- Data in the tool was updated to reflect USGS information as of 02/04/2022
- Added the ability to view segment trend lines in breakpoint analysis
- Redesigned breakpoint table to show breakpoints by segments and their associated regression statistics
- Added in customization options for time series modeling parameters
- Added in confidence intervals in time series model visualization for fitted time models
- Change the seasonal decomposition file output format to include decomposed elements as column headers
- Integrated upload data page to include preloaded USGS data
- Integrated method explorer into the NSD main tab
- Added in gage statistics and gage map within the upload data panel
- Incorporated gage data table into the NSD main tab
- Change breakpoint analysis missing data handling to match NSD tab and concatenate segments together if there is missing data
- Allow seasonal cycle analysis to be performed for data aggregated to the monthly level

Version 2.1 Change Log

- Data in the tool was updated to reflect USGS information as of 08/01/2023
- Added preloaded monthly streamflow/gage height data



- Added option to allow the user to directly analyze seasonally decomposed data components
- Updated the downloadable nonstationarity report that showcases the three graphs from the NSD tab
- Updated the user guide to reflect the literature review conducted for understanding seasonality and its effects on nonstationarity detection
- As a result of a literature review and experimentation on tidal gauge data, added in autocorrelation metrics to the summary statistics page and a detrending preprocessing option for user uploaded data
- Improved aesthetics for the nonstationarity graphics