

A SEQUENTIAL METHOD FOR DETECTING REGIME SHIFTS IN THE MEAN AND VARIANCE

Sergei N. Rodionov

Joint Institute for the Study of the Atmosphere and Ocean
University of Washington, Seattle, Washington

Introduction

In interpreting long-term variations in climatic and biological records, a concept of “regimes” and “regime shifts” has become very popular in recent decades. This concept received a strong impetus after a step-like change in the global climate system in the late 1970s, although a realization of the importance of that event came more slowly. A number of methods have been developed to detect regime shifts or change points in time series (see an overview of these methods by Rodionov in this volume). The overwhelming majority of these methods are designed to find shifts in the mean, and only a few can do this for the variance. Changes in the variance of climatic parameters may have similar or even greater impact on marine ecosystems than changes in the mean. As climate changes due to natural causes or human impact, it may be changes in the frequency of hazards or extreme events that pose the most immitate danger than the increase in the mean surface level temperature referred to as „global warming”.

Most of the reviewed methods have one common drawback: their performance drastically deteriorates if change points are too close to the ends of time series. A possible solution to this problem lies in the use a sequential data processing technique. In sequential analysis the number of observations is not fixed. Instead, observations come in sequence. For each new observation a test is performed to determine the validity of the null hypothesis H_0 (the existence of a regime shift in this case). There are three possible outcomes of the test: accept H_0 , reject H_0 , or keep testing. Recently Rodionov (2004) introduced a sequential method for detecting regime shifts in the mean that was tested on a set of indices describing the Bering Sea ecosystem (Rodionov and Overland, 2005). Below is the functional description of the method and its extension for detecting shifts in the variance.

Shift in the Mean

The method is based on the sequential application of the Student’s t-test, which is used here in the spirit of exploratory, rather than confirmatory, data analysis. Let $x_1, x_2, \dots, x_i, \dots$ be a time-series with new data arriving regularly. When a new observation arrives, a check is

performed to determine whether it represents a statistically significant deviation from the mean value of the „current” regime (\bar{x}_{cur}). According to the t-test, the difference between and the mean value of the new regime (\bar{x}_{new}) to be statistically significant at the level p should satisfy the conditions

$$diff = |\bar{x}_{new} - \bar{x}_{cur}| = t\sqrt{2\bar{s}_l^2/l}$$

where t is the value of the t -distribution with $2l - 2$ degrees of freedom at the given probability level p . It is assumed here that the variances for both regimes are the same and equal to the average variance for running l -year intervals in the time series $\{x_i\}$. It means that $diff$ remains constant for the entire session with the given time series.

At the „current” time t_{cur} , the mean value of the new regime \bar{x}_{new} is unknown, but it is known that it should be equal or greater than the critical level $\bar{x}_{crit}^{\uparrow}$, if the shift is upward, or equal or less than $\bar{x}_{crit}^{\downarrow}$, if the shift is downward, where

$$\bar{x}_{crit}^{\uparrow} = \bar{x}_{cur} + diff,$$

$$\bar{x}_{crit}^{\downarrow} = \bar{x}_{cur} - diff.$$

If the current value x_{cur} is greater than $\bar{x}_{crit}^{\uparrow}$ or less than $\bar{x}_{crit}^{\downarrow}$, the time t_{cur} is marked as a potential change point c , and subsequent data are used to reject or accept this hypothesis. The testing consists of calculating the so-called regime shift index (RSI) that represents a cumulative sum of normalized anomalies relative to the critical level \bar{x}_{crit} :

$$RSI = \frac{1}{l\bar{s}_l} \sum_{i=t_{cur}}^m (x_i - \bar{x}_{crit}), \quad m = t_{cur}, t_{cur} + 1, \dots, t_{cur} + l - 1.$$

If at any time during the testing period from t_{cur} to $t_{cur} + l - 1$ the index turns negative, in the case of $\bar{x}_{crit} = \bar{x}_{crit}^{\uparrow}$, or positive, in the case of $\bar{x}_{crit} = \bar{x}_{crit}^{\downarrow}$, the null hypothesis about the existence of a shift in the mean at time t_{cur} is rejected, and the value x_{cur} is included in the „current” regime. Otherwise, the time t_{cur} is declared a change point c .

Shift in the Variance

The procedure for detecting regime shifts in the variance is similar to the one for the mean, except that it is based on the F-test instead of the t-test. It is assumed that the mean value of the time series is zero, that is, we work with the residuals $\{z_i\}$ after shifts in the mean are removed from the original time series $\{x_i\}$. The F-test consists of comparing the ratio of the sample variances for two regimes with the critical value F_{crit} :

$$F = \frac{S_{cur}^2}{S_{new}^2} > F_{crit}.$$

Here F_{crit} is the value of the F-distribution with \hat{l}_1 and ν_2 degrees of freedom (where $\hat{l}_1 = \nu_2 = l - 1$) and a significance level p (two-tailed test):

$$F_{crit} = F(p/2, \hat{l}_1, \nu_2).$$

The variance s_{cur}^2 is the sum of squares of z_i , where i spans from the previous shift point in the variance (which is the first point of the “current” regime) to $i = t_{cur} - 1$. At the “current” time t_{cur} , the variance s_{new}^2 is unknown. For the new regime to be statistically different from the current regime, the variance s_{new}^2 should be equal or greater than the critical variance $s_{crit}^{2\uparrow}$, if the variance is increasing, or equal or less than $s_{crit}^{2\downarrow}$, if the variance is decreasing, where

$$s_{crit}^{2\uparrow} = s_{cur}^2 F_{crit},$$

$$s_{crit}^{2\downarrow} = s_{cur}^2 / F_{crit}.$$

If at the time t_{cur} the current value z_{cur} satisfies one of the following conditions, $z_{cur}^2 > s_{crit}^{2\uparrow}$ or $z_{cur}^2 < s_{crit}^{2\downarrow}$, this time is marked as a potential shift point, and subsequent values z_{cur+1} , z_{cur+2} ... are used to verify this hypothesis. The verification is based on the residual sum of squares index (RSSI) defined as

$$RSSI = \frac{1}{l} \sum_{i=t_{cur}}^m (z_i^2 - s_{crit}^2), \quad m = t_{cur}, t_{cur} + 1, \dots, t_{cur} + l - 1.$$

The decision rule is similar to the one for shifts in the mean: If at any time during the testing period from t_{cur} to $t_{cur} + l - 1$ the index turns negative, in the case of $s_{crit}^2 = s_{crit}^{2\uparrow}$, or positive, in the case of $s_{crit}^2 = s_{crit}^{2\downarrow}$, the null hypothesis about the existence of a shift in the variance at time t_{cur} is rejected, and the value z_{cur} is included in the “current” regime. Otherwise, the time t_{cur} is declared a change point c .

An Example

The above procedures were coded using Visual Basic for Applications (VBA) in the form of an Excel Add-In. It is available for download from the Bering Climate web site (www.BeringClimate.noaa.gov). The website also contains detailed instructions on how to install and use the add-in.

Using this add-in, I first generated a time series of annual values from 1901 to 1960, which consisted of three 20-yr segments. The first segment is a realization of a normal (Gaussian) process with the zero mean value and the variance of 0.5, $N(0, 0.5)$, the second segment is $N(2, 4)$, and the third segment is again $N(0, 0.5)$. This time series is presented in Fig. 1a.

The first step of the analysis is to remove shifts in the mean. To do so, the program is run using the following parameters: cutoff length $l = 10$ years and probability level $p = 0.1$. The regime shifts were detected at 1921 and 1939. Some variations in l and p (for example, setting $l = 20$ and $p = 0.05$) produce the same results. In terms of the t-criterion, the significance level for the 1939 shift is practically the same as the one for the theoretical shift in 1941. The program tends to pick the first change point that satisfies the given conditions. Only when the probability level is reduced to $p = 0.01$, the change point is detected at 1941.

The test for shifts in the variance is performed on the residuals (Fig. 1b), after the stepwise trend (gray line in Fig. 1a) is removed, using the same l and p values as for the mean. Positive RSSI values were obtained for 1923 and 1941 (Fig. 1c). Based on the F-criterion, the significance level for the shift in 1923 is lower than the one for the theoretical shift in 1921.

Summary of the Method's Features

- The method is fully automatic and capable of detecting multiple change-points in a time series. It does not require an a priori hypothesis on the timing of regime shifts, which eliminates the problem of data-dredging that arises in testing for change occurring at a specified time (Epstein, 1982).

- It can be tuned to detect the regimes of certain time scales and magnitudes. The time scale to be detected is controlled primarily by the cut-off length l . As the cut-off length is reduced, the time scale of regimes detected becomes shorter. Both the cut-off length l and probability level p affect the statistically significant difference between regimes, and hence the magnitude of the shifts to be detected. Note that the value of p set for the time series is the maximum significance level at which regimes shifts can be detected. Actual significance levels for the differences between regimes (which are also calculated) are usually less than p .

- It can handle the incoming data regardless whether they are presented in the form of anomalies or absolute values. This eliminates the necessity to select the base period to calculate anomalies, which is a source of ambiguity that affects the timing and scale of the regimes.

- It can be applied easily to a large set of variables. It is important to note that there is no need to reverse the sign of some time-series to ensure that all shifts occur in the same direction. This problem was experienced by Hare and Mantua (2000) in their analysis of 100 physical and biological time-series in the North Pacific. Rudnick and Davis (2003) demonstrated that the procedure of sign reversal artificially enhances the chance of identifying existing shifts and may even lead to spurious shifts being identified.

- It is quite robust in relation to a linear trend in the time series. If a trend is present in a time series, it may create a serious problem because it is easy to falsely identify as a shift point the center of this time series.

- Perhaps the most important feature of the proposed method may be its ability to detect a regime shift relatively early and then monitor how its magnitude changes over time.

Concluding Remarks

The method assumes that each data point is independent of the other measurements, so that there is no serial correlation (autocorrelation). Although the method is quite robust to the assumption of data independence, the existence of a strong autocorrelation in the time series can lead to an increased number of incorrectly identified regime shifts („false alarms”). Two approaches are possible to overcome this problem. First, the test formulas can be modified to take into account the existence of autocorrelation. The overall effect would be equivalent to increasing the probability level p . The second approach is to perform the so-called „prewhitening,”

a procedure that removes the red noise component (caused by the autocorrelation) from the time series. A new version of the computer program that includes the „prewhitening” procedure is expected to be posted on the Bering Climate website by the end of 2005.

References

- Epstein, E. S. (1982). Detecting climate change, *J. Appl. Meteorol.*, 21, 1172.
- Hare, S. R. and N. J. Mantua. (2000). Empirical evidence for North Pacific regime shifts in 1977 and 1989, *Progr. Oceanog.*, 47, 103-146.
- Rodionov, S. (2004): A sequential algorithm for testing climate regime shifts, *Geophys. Res. Lett.*, 31, L09204, doi:10.1029/2004GL019448.
- Rodionov, S. and J. E. Overland. (2005). Application of a sequential regime shift detection method to the Bering Sea ecosystem, *ICES Journal of Marine Science*, 62, 328-332.
- Rudnick, D. I. and R. E. Davis. (2003). Red noise and regime shifts, *Deep-Sea Research*, 50, 691-699.




Fig. 1. a) A synthetic time series consisting of three segments of normally distributed random numbers with the following mean values and variances: 1) 0, 0.5, 2) 2, 4, and 3) 0, 0.5. Gray line is the stepwise trend showing regime shifts in the mean detected by the sequential method; b) the same time series after removing the stepwise trend; and c) RSSI showing regime shifts in the variance.