

Detection of signals by Monte Carlo singular spectrum analysis: multiple testing*

NINA GOLYANDINA

Detection of a signal in a noisy time series using Monte Carlo singular spectrum analysis (MC-SSA) is studied from the statistical viewpoint. The MC-SSA test consists of simultaneous testing of several hypotheses related to the presence of different frequencies. The multiple MC-SSA test procedure is constructed to control the family-wise error rate. The technique to control both the type I and the type II errors and also to compare criteria is proposed to study several versions of MC-SSA.

AMS 2000 SUBJECT CLASSIFICATIONS: Primary 62G10, 94A12; secondary 60G35, 37M10.

KEYWORDS AND PHRASES: Singular spectrum analysis, Time series, Signal detection, Multiple testing, Family-wise error rate.

1. INTRODUCTION

Singular spectrum analysis (SSA) solves a wide range of problems of time series analysis and image processing (see [6] for examples and references). Here we consider a specific problem of detection of a signal (e.g., a periodic component) in a noisy time series.

The common scheme of the decomposition stage of SSA consists of the construction of a so-called trajectory matrix \mathbf{X} from the initial object \mathbf{X} and an expansion of this trajectory matrix into a sum of elementary components, which are ordered by their contribution. In the basic version of SSA, this expansion is performed using the singular value decomposition (SVD). The left singular vectors of the SVD are eigenvectors of the matrix $\mathbf{X}\mathbf{X}^T$, which is an estimate of the autocovariance matrix of the time series if its mean equals zero or the time series is centred. The reconstruction stage starts with the identification of elementary components corresponding to a component under interest (e.g. a signal) and then the reconstruction of this time series component through a grouping of the identified components is performed.

Generally, SSA is a model-free method. However, a considerable part of the SSA theory is devoted to the extraction of time series components, which are governed (maybe,

approximately) by a linear recurrence relation [6]; in particular, a sum of sine waves with slowly-varying amplitudes belong to this class of signals. SSA extracts periodic components well; however, it is well-known that the extracted components can be spurious, since they can be produced by noise. In a sense, this is a payment for the nonparametric nature of the method. If we want to apply the statistical approach for judging, a model should be assumed. Usually, the question about the existence of a signal in the time series is formulated as testing the hypothesis that the series is a stochastic process. The criterion should be powerful against the alternative hypothesis, which states the existence of a non-random (e.g. periodic) component. There are a lot of statistical criteria for testing these hypotheses for different classes of stochastic processes (see e.g. [17] for red noise or [13] for an Ornstein-Uhlenbeck state space, which is a continuous analogue of the first-order autoregressive processes AR(1)). We consider the construction of such a criterion in the framework of SSA, since SSA is able to reconstruct the detected time series component.

It follows from the properties of SSA (see e.g. the description of the relation between the spectral density and the eigenvectors in [7, Sect. 6.4.3], where the results from [9] in terms of SSA are discussed) that a natural assumption for the noise model is that noise is red (the AR(1) with a positive coefficient). This is because the spectral density of red noise is monotonic; therefore, the eigenvectors are similar to sinusoids and can be connected to different frequencies. Recall that $\boldsymbol{\xi} = (\xi_1, \dots, \xi_N)$ is red noise with parameters φ and δ if $\xi_n = \varphi\xi_{n-1} + \delta\epsilon_n$, where $0 < \varphi < 1$, ϵ_n is white Gaussian noise with mean 0 and variance 1 and ξ_1 has a normal distribution with mean zero and variance $\delta^2/(1 - \varphi^2)$. Moreover, in climatology, the common model is a weak signal (if any) in red noise. Both considerations inspired the creation of the Monte Carlo SSA method (MC-SSA). MC-SSA was proposed in [2] and later was considered in many papers ([1], [4], [8], [10], [14], [15], among others).

However, the terminology in these papers differs from the standard statistical terminology and therefore it is very important to bridge the gap between the applied and statistical approaches. Moreover, this can help to avoid wrong conclusions in real-life applications. Investigation of criterion properties (and comparison with other tests) will be performed by the estimation of the type I and type II errors.

Let us briefly describe how the basic MC-SSA test is constructed. At Decomposition step of SSA, each decom-

*The reported study was funded by RFBR, project number 20-01-00067.

position component (related to the eigenvectors of the sample autocovariance matrix, in the basic version) can be put into correspondence with a frequency, since the eigenvectors of the autocovariance matrix of red noise, which is a Kac–Murdock–Szegő Toeplitz matrix, are sine waves with almost equidistant frequencies [16]. The eigenvalues are equal to the total squared norm of projection of the trajectory matrix columns on the eigenvectors and therefore reflect contributions of the decomposition components (and therefore of the corresponding frequencies) into the time series. The original idea of MC-SSA is to estimate the parameters of red noise and apply the bootstrap simulations to construct prediction intervals for eigenvalues in the case when there is no signal. If an eigenvalue of the time series is beyond the constructed prediction intervals, the corresponding eigenvector frequency is considered significant. Moreover, it is possible to reconstruct the detected signal. In the modifications of MC-SSA, the choice of vectors for projection can vary.

It is clear from the method description that there is the problem of multiple testing when the probability of the false detection of a periodic component for one of the considered frequencies (family-wise error rate) is unknown and is much larger than the given significance level (single-test error rate). This problem is discussed in different papers devoted to MC-SSA. The theoretical approach to multiple testing, which we propose in this paper, allows constructing the multivariate criterion with the given family-wise error rate.

Novelty. The novelty of the paper is applying the statistical approach to the signal-detection problem in the framework of Monte Carlo SSA to control the type I error and estimate the type II error. For simultaneous testing of the presence of multiple frequencies, a multiple version of MC-SSA is proposed to control the family-wise error rate. Based on the numerical study, we discuss pluses and minuses of several versions of MC-SSA.

Structure. In Section 2, we briefly describe the statistical approach to hypothesis testing. In Section 3, versions of the MC-SSA algorithms are presented, including the proposed multiple testing algorithms. Section 4 is devoted to the investigation of the MC-SSA approach using the statistical approach. All numerical studies were performed in R, with the use of the RSSA package [11].

2. STATISTICAL APPROACH TO HYPOTHESIS TESTING

In papers starting from [2], the method MC-SSA is described mostly as a method for applied problems and therefore the way of description is not conventional for statisticians. Therefore, let us start with a statistical approach to the problem.

Let the null hypothesis be that the time series is a pure stationary stochastic process. In the considered context, it can be white or red noise. Sometimes, one says that the

presence of a signal in noise is tested, whereas the null hypothesis is formulated as the hypothesis about the absence of a signal in noise. Consider a criterion, which determines if the null hypothesis is rejected or is not rejected. If the null hypothesis is rejected at the given significance level α , then one can claim the presence of a signal (more precisely, a deviation from the null hypothesis). The probability to reject the null hypothesis if it is true is called type I error (α_I). If a criterion is correct, then the type I error is equal to the given significance level (or at least not larger than α). Different criteria differ by the power against an alternative hypothesis. The power is the probability to reject the null hypothesis if the alternative hypothesis is true. The alternative hypothesis that the time series contains a periodic component is important in practice; therefore, we will consider criteria, which are powerful against such hypotheses.

If we want to make a false discovery rarely, then we choose a small significance level α and this guarantees that we will reject the true null hypothesis with probability not larger than α . However, we can not choose a very small significance level, since the test power decreases as α decreases.

Note that it is not permitted to consider a criterion if its type I error exceeds the given significance level (a *liberal* criterion). Therefore, before a comparison of criteria by power, one should be sure that the type I error lies in the given range. If the type I error is less than the significance level (a *conservative* criterion), this is admissible; however, this means that this criterion can be improved, that is, the power can be increased by a correction to obtain the type I error closer to the significance level.

A useful characteristic of a criterion is the possibility to interpret the difference from the null hypothesis if this hypothesis is rejected.

2.1 Bootstrapping

Most of the criteria have the following form. A constructed test statistic measures the difference between data and the null hypothesis in some way. There is a threshold such that if the test statistic is larger than the threshold, the null hypothesis is rejected. Certainly, this threshold depends on the significance level α . It is not uncommon that this threshold can not be obtained theoretically. Then simulations are used. Surrogate data are simulated according to the null hypothesis and the test statistic is calculated many times (G) to determine the threshold; this approach is widely used, see e.g. [12]. Consequently, the threshold, for which the null hypothesis is rejected approximately αG times, is found. The estimated threshold is used for testing the hypothesis for real-life data. The surrogate data should be obtained exactly in the same way as the test statistic was generated for the original data. The described approach can be called Monte-Carlo one. This approach helps to construct the criterion with the type I error tending to α as G tends to infinity.

However, the Monte Carlo approach can be applied if the null hypothesis fully determines the data. For example, the null hypothesis states that the time series is red noise with variance δ^2 and coefficient φ , where δ^2 and φ are known numbers.

Unfortunately, this is not the case in practice. Therefore, the so-called bootstrapping is used (“pull yourself up by your bootstraps”). If δ^2 and φ are unknown, then they are estimated with the help of the real-life data under study and then the surrogate data are produced by simulations with the estimated parameters. Since the estimated parameters differ from the true (unknown) parameters, then the type I error can generally be far from α and the test can become liberal or conservative. Thus, Monte-Carlo SSA (this is the name of a family of concrete algorithms) is a kind of testing with bootstrapping.

2.2 Estimation of type I error and power

The above considerations about the relation between the type I error, the significance level α and the level of power can not be applied in practice, since the type I error and the power are unknown. If something is unknown theoretically, then simulation helps again.

Let a criterion be constructed to make the decision (reject or not reject) for a given significance level α . It can be constructed theoretically or with bootstrapping/simulations within, it does not matter. Let the sample data with given parameters according to the null hypothesis (this is the Monte-Carlo approach) be simulated many times (M). Then the proportion of cases with the rejection of the null hypothesis is an estimate of the type I error.

To estimate power, we should set an alternative hypothesis. There is freedom in the choice. The common rule is to include into the alternative hypothesis the case, which is important, that is, the case that should be distinguished from the null hypothesis. For example, the alternative can state that the time series is a sine wave with a given frequency, amplitude and phase corrupted by noise with the same parameters as were considered for the null hypothesis.

For power estimation, the procedure similar to that for the type-I error is fulfilled. We simulate surrogate data with the given parameters according to the alternative hypothesis many times (M). Then the proportion of cases with the rejection of the null hypothesis is the estimate of the power.

2.3 Prediction intervals for testing hypotheses

Let a test statistic t provide an interpreted characteristic of data (e.g. a contribution of a frequency to the observed time series). Thus, the question about the presence of a signal can be reformulated as “can this value of t be caused by the noise component only?”.

The answer to the question can be obtained in the standard way. Since the contribution t is random, there is a

prediction interval for it. The prediction interval can be constructed by simulation. If we generated a sample of possible contributions, then the 95% prediction interval is the interval between the 2.5% and 97.5% quantiles of this sample. In statistical terminology, this interval is not called confidence, since confidence intervals are constructed for unknown (non-random) parameters and their length would tend to zero as the number of simulations (sample size) tends to infinity.

γ -Prediction intervals serve for testing hypotheses with a significance level α for $\alpha = 1 - \gamma$. If the observed value of t does not belong to the prediction interval, the null hypothesis is rejected. It is convenient to depict, say, 95%-prediction intervals for t to visualize the hypothesis testing with the significance level 5%.

2.4 One-tailed and two-tailed criteria

We mentioned that a criterion consists of a test statistic t and a threshold t_0 . The use of this threshold can be different. Moreover, the threshold can consist of two numbers, t_1 and t_2 . For example, the null hypothesis can be rejected if the test statistic is larger than t_1 or smaller than t_2 (two-tailed test) or if $t > t_0$ (one-tailed test). The choice of the criterion type depends on the alternative hypothesis, since we want to increase the power against the chosen alternative.

If we want to detect both cases, when the contribution of a frequency is either larger or smaller than that for pure noise, then we choose a two-tailed test. If we want to detect only the excess of the frequency contribution, we choose a one-tailed test. This approach can be expressed in terms of prediction intervals. If we are interested to find the frequency with contribution larger than that of noise, then the one-sided prediction interval has the form $[0, t_0]$ (in the general case, $[-\infty, t_0]$; however, in our case the test statistic is non-negative). In the two-tailed case, the two-sided prediction interval is $[t_1, t_2]$.

2.5 Multiple testing

The problem of multiple testing is well-known. The approach to the statistical testing described above is applicable for single tests only, since the probability of false discovery is controlled for each individual test only.

If we test several tests (m) simultaneously, we are interested in the so-called family-wise error rate (FWER). FWER is the probability of false discovery in at least one of m tests. This probability can be vastly larger than the chosen small α . Thus, we should not use a set of single tests with α if we want to control FWER. Ideally, we should construct one multivariate test instead of several single tests. If this is technically hard, the Bonferroni correction is used (performing single tests with significance level α/m); this trick controls the FWER not larger than α (usually, this does the family-wise type I error considerably smaller than α ; that is, the multiple testing is conservative and therefore decreases the test power). If the single tests are independent,

the Šidák correction (performing single tests with significance level $1 - (1 - \alpha)^{1/m}$) provides the exact test; however, in the general case, testing with the Šidák correction can be liberal and therefore the Šidák correction is not applied.

3. MONTE CARLO SSA

3.1 Singular spectrum analysis

Let us shortly describe the scheme of SSA to introduce notation used further (see e.g. [6] for details).

Denote by $\mathbf{X} = (x_1, \dots, x_N)$ a time series of length N and by L , $1 < L < N$, a window length. The *trajectory matrix* $\mathbf{X} = \mathcal{T}(\mathbf{X})$ is determined as $\mathbf{X} = [X_1 : \dots : X_K]$, where $K = N - L + 1$ and $X_i = (x_i, \dots, x_{i+L-1})^T \in \mathbb{R}^L$ are lagged vectors.

The next step is the SVD expansion $\mathbf{X} = \sum_{i=1}^d \sqrt{\lambda_i} U_i V_i^T = \sum_{i=1}^d U_i (\mathbf{X}^T U_i)^T$, where $\{U_i\}_{i=1}^d$ is the orthonormal system of the eigenvectors of the matrix $\mathbf{X}\mathbf{X}^T$, $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d > 0$ are the corresponding non-zero eigenvalues. This is Basic SSA (or BK version); sometimes Toeplitz SSA (VG version) designed for stationary time series is considered.

The components of the obtained matrix decomposition are reasonably grouped and each grouped matrix is transferred to a time series. Thus, the result of SSA is a time series decomposition.

Although SSA is a model-free method, there is a model that fits it. This is a class of signals, which are approximated by a sum of products of polynomials, exponentials and sine waves. In particular, a sum of sine waves is perfect for SSA. SSA can extract sine waves with different frequencies if these frequencies are not too close and can separate a sum of sinusoids from noise. If a sine wave series component can be extracted from the time series (we say that it is separated from the residual by SSA), then the SVD decomposition contains two eigenvectors U_i and U_{i+1} , which have the same dominant frequency as the original sine wave.

By the properties of the SVD, $\lambda_i = \|\Pi_{\{U_i\}} \mathbf{X}\|^2 = \|\mathbf{X}^T U_i\|^2 = \sum_{j=1}^K (X_j^T U_i)^2$ can be interpreted as the total squared norm of projections of lagged vectors to $\text{span}(U_i)$.

3.2 General comments

On the one hand, Monte Carlo SSA is a well-developed method. In [2], the foundation of the method is thoroughly described. On the other hand, many questions are still under investigation; they are, among others, the best way of estimating the noise parameters and taking into consideration the presence of the nuisance signal. In this paper, we consider a most simple version of Monte Carlo SSA to concentrate on the problem of multiple testing.

In some cases, we will assume that the noise parameters are known. Also, we will show how can the criterion errors change if the parameters are estimated.

There is a natural modification of Monte Carlo SSA, which does not use the SSA decomposition step except for

the construction of the trajectory matrix (a part of the study will be performed for such a modification). However, the relation with SSA is essential, since the detected signal can be reconstructed through SSA if we keep the connection with the eigenvectors in SSA.

3.3 Single test

Let $\boldsymbol{\xi} = (\xi_1, \dots, \xi_N)$ be a red noise with parameters φ and δ . Denote by L the window length and by $\boldsymbol{\Xi} = \mathcal{T}(\boldsymbol{\xi})$ the trajectory matrix of $\boldsymbol{\xi}$. Let a vector $W \in \mathbb{R}^L$, $\|W\| = 1$, be given. If we are interested in a frequency contribution, then W can be a sine wave with a given frequency. The total squared norm of the projection of the columns of $\boldsymbol{\Xi}$ to the vector W is calculated as

$$p = \|\boldsymbol{\Xi}^T W\|^2 = W^T (\boldsymbol{\Xi} \boldsymbol{\Xi}^T) W.$$

The null hypothesis states that the observed time series \mathbf{X} is a realization of $\boldsymbol{\xi}$ with some parameters φ and δ . Denote $\hat{p} = \|\mathbf{X}^T W\|^2$. If W is an eigenvector of $\mathbf{X}\mathbf{X}^T$, then \hat{p} is the corresponding eigenvalue. Note that for a sinusoidal W , \hat{p} negligibly depends on the phase of this sinusoid, since for large $K = N - L + 1$ the lagged vectors consist of many shifts.

Let φ and δ be known. Under some assumptions, the distribution of p can be calculated theoretically. Then the prediction interval with confidence level γ is calculated as the interval between $(1 - \gamma)/2$ - and $(1 + \gamma)/2$ - quantiles for the two-tailed test and between zero and γ -quantile for the one-tailed test (upper-tailed test). In both cases, \hat{p} belongs to the constructed predicted interval with probability γ .

If the theoretical distribution is unknown, then these quantiles can be calculated by simulation of G samples $\boldsymbol{\xi}^i$ of the random vector $\boldsymbol{\xi}$ and the use of empirical (sample) quantiles for the obtained sample $p_i = \|\boldsymbol{\Xi}_i^T W\|^2$, $i = 1, \dots, G$. The probability that \hat{p} belongs to the empirical (Monte Carlo) prediction interval tends to γ as G tends to infinity.

Recall that the significance level α is equal to $1 - \gamma$ and therefore one can say that the probability of the type I error α_I tends to α .

For both theoretical considerations and simulations, the values of the parameters φ and δ are used. Here we do not discuss the estimation of noise parameters in the presence of a signal. Note that the proposed approach to hypothesis testing is valid for various modifications of the MC-SSA technique.

Remark 1. It is important to note that W may be produced by the time series itself. However, as was discussed in [2], then the surrogate data should be projected also to the vectors produced by them. There is a temptation to consider the version when the projection vector is produced by the observed time series and the test is constructed by projection to this vector. However, then the type-I error is not controlled and the test is liberal.

3.4 Choice of vectors for projection

In practice, we do not know the frequency of a possible periodic signal component. Therefore, the approach is to consider many vectors for projection, which correspond to a set of frequencies. For example, one can take a set of sine waves $W_1, \dots, W_H \in \mathbb{R}^L$ with equidistant frequencies from some frequency interval $[\omega_1, \omega_2] \subset (0, 0.5)$. To obtain slightly dependent tests, the number of vectors should not exceed their dimension L .

The other choice is to take the set of eigenvectors produced by SSA (this is a common case in MC-SSA) with consideration to Remark 1.

The compromised version suggested in [2] is W_1, \dots, W_H to be eigenvectors of the general-population covariance matrix. That is, the sample covariance matrix is $\mathbf{X}\mathbf{X}^T$, where \mathbf{X} is the L -trajectory matrix, whereas the general-population correlation matrix of red noise has the (i, j) -th term $\phi^{|i-j|}$. They are close; however, the difference from the viewpoint of type-I error can be drastic. If one wants to include to the set of projection vectors the sine wave vector with a specific frequency ω , then the eigenvectors of the matrix with the (i, j) -th term $\phi^{|i-j|} + C \cos(2\pi\omega|i-j|)$ can be considered.

Thus, we will consider two choices:

1. W_1, \dots, W_H are the eigenvectors of the matrix

$$(1) \quad \{\phi^{|i-j|} + C \cos(2\pi\omega|i-j|)\}_{i,j=1}^L.$$

2. W_1, \dots, W_H are the eigenvectors of the matrix $\mathbf{X}\mathbf{X}^T$; $H \leq \min(L, N - L + 1)$.

Since we study the version of Monte Carlo SSA, where the projections of surrogate data are performed on the fixed vectors, the second version is generally wrong. We will consider an example to show this.

3.5 Multiple testing

In Monte Carlo SSA, the prediction intervals are constructed for the contribution of each projection vector independently. Let W_1, \dots, W_H be a set of projection vectors. Denote

$$\hat{p}_k = \|\mathbf{X}^T W_k\|^2, \quad k = 1, \dots, H.$$

For each vector W_k , the sample of squared projection norms is constructed: $P_k = (p_{k1}, \dots, p_{kG})^T$, where p_{ki} is calculated as

$$(2) \quad p_{ki} = \|\Xi_i^T W_k\|^2, \quad i = 1, \dots, G;$$

here $\Xi_i = \mathcal{T}(\xi^i)$ is the trajectory matrix of the i th sample of red noise $\xi^i = (\xi_1^i, \dots, \xi_N^i)$ with given parameters.

We can construct single prediction intervals for the contribution of each vector W_k as it is described in Section 3.3. The problem of multiple testing (the problem of FWER, which can be much larger than the given significance level α) can be solved by means of the Bonferroni correction. However, if the Bonferroni correction is used, the test becomes

conservative (FWER is less than α). To obtain an exact test, an approach similar to that of Tukey's HSD applied to post hoc comparisons in ANOVA can be considered. That is, we can construct a test, which is based on the distribution of the maximum of the standardized contributions p_k . If this test rejects the null hypothesis, then all frequencies, which correspond to the projection vectors with the contribution lying outside the corrected prediction intervals, are considered significant. Thus, we can talk about the significance of a frequency if put into correspondence projection vectors and frequencies. Note that for this approach, FWER is equal to α .

Let us describe the algorithm of constructing the prediction intervals with correction for multiple testing.

The first version (Algorithms 1 and 2) is straightforward; this is the single test for each vector W_k with Bonferroni correction, that is, the significance level α/H is taken instead of α .

Algorithm 1 (Single one-tailed test with Bonferroni correction)

1. For each k , $k = 1, \dots, H$, calculate the test statistic \hat{p}_k and the sample $P_k = \{p_{ki}\}_{i=1}^G$, see (2).
2. Find q_k as the sample $(1 - \alpha/H)$ -quantile of P_k .
3. The null hypothesis, which states that the time series is pure red noise, is not rejected if for each k the inequality $\hat{p}_k < q_k$ is valid; otherwise, the null hypothesis is rejected and a signal is detected.
4. If H_0 is rejected, then post-hoc testing can be performed: the contribution of W_k (and of the corresponding frequency) is significant if \hat{p}_k exceeds q_k . Thus, $[0, q_k]$ are considered as the corrected prediction intervals, $k = 1, \dots, H$.

Algorithm 2 (Single two-tailed test with Bonferroni correction)

1. For each k , $k = 1, \dots, H$, calculate the test statistic \hat{p}_k and the sample $P_k = \{p_{ki}\}_{i=1}^G$, see (2).
2. Find q_k^{low} and q_k^{up} as the sample $(0.5\alpha/H)$ - and $(1 - 0.5\alpha/H)$ -quantiles correspondingly.
3. The null hypothesis, which states that the time series is pure red noise, is not rejected if for each k $q_k^{\text{low}} < \hat{p}_k < q_k^{\text{up}}$; otherwise, the null hypothesis is rejected and a signal is detected.
4. If H_0 is rejected, then post-hoc testing can be performed: the contribution of W_k (and of the corresponding frequency) is significant if \hat{p}_k does not belong $[q_k^{\text{low}}, q_k^{\text{up}}]$. Thus, $[q_k^{\text{low}}, q_k^{\text{up}}]$ are considered as the corrected prediction intervals, $k = 1, \dots, H$.

Let us extend the multiple-testing approach, which was considered in [3] in the framework of Monte Carlo SSA. The approach is based on the distribution of $\max_{1 \leq k \leq H} (p_{ki} - \mu_k) / \sigma_k$, where μ_k and σ_k are mean and standard deviation

of P_k , $k = 1, \dots, H$. Here σ_k reflects the size of the k -th prediction intervals. Normalization by σ_k keeps the same difference between the intervals sizes as in the single tests. By construction, the size of a prediction interval is related to the test power against the existence of the signal at the corresponding frequency. In Algorithms 3 and 4, we consider $\tilde{\sigma}_k = w_k \sigma_k$, where w_k is a weight; an increase of the weight corresponds to more expectation of a signal at the frequency matched to W_k . Some weights w_k can be zero; this means that the corresponding projection vectors do not participate in the testing and the test is disabled to detect the signal at the corresponding frequencies. Thus, it is the same, to set zero weights outside the frequency range (and unit weights within it) or just consider the subset of the projection vectors matched to the frequency interval. However, the approach with weights is much more flexible and allows to take arbitrary weights between 0 and 1.

Algorithm 3 (Weighted multiple one-tailed test)

1. For each k , $k = 1, \dots, H$, calculate the test statistic \hat{p}_k , the sample $P_k = \{p_{ki}\}_{i=1}^G$, see (2), and calculate its mean μ_k and standard deviation σ_k .

2. Calculate $\boldsymbol{\eta} = (\eta_1, \dots, \eta_G)$, where

$$\eta_i = \max_{1 \leq k \leq H} (p_{ki} - \mu_k) / \tilde{\sigma}_k, \quad \tilde{\sigma}_k = w_k \sigma_k, \quad i = 1, \dots, G.$$

3. Find q as the sample $(1 - \alpha)$ -quantile of $\boldsymbol{\eta}$.

4. The null hypothesis, which states that the time series is pure red noise, is not rejected if

$$\max_{1 \leq k \leq H} (\hat{p}_k - \mu_k) / \tilde{\sigma}_k < q;$$

otherwise, the null hypothesis is rejected and a signal is detected.

5. If H_0 is rejected, then post-hoc testing can be performed: the contribution of W_k (and of the corresponding frequency) is significant if \hat{p}_k exceeds $\mu_k + qw_k \sigma_k$. Thus, $[0, \mu_k + qw_k \sigma_k]$ are considered as the corrected prediction intervals, $k = 1, \dots, H$.

Algorithm 4 (Weighted multiple two-tailed test)

1. For each k , $k = 1, \dots, H$, calculate the test statistic \hat{p}_k , the sample $P_k = \{p_{ki}\}_{i=1}^G$, see (2), and calculate its mean μ_k and standard deviation σ_k .

2. Calculate $\boldsymbol{\eta} = (\eta_1, \dots, \eta_G)$, where

$$\eta_i = \max_{1 \leq k \leq H} |p_{ki} - \mu_k| / \tilde{\sigma}_k, \quad \tilde{\sigma}_k = w_k \sigma_k, \quad i = 1, \dots, G.$$

3. Find q as the sample $(1 - \alpha)$ -quantile of $\boldsymbol{\eta}$.

4. The null hypothesis, which states that the time series is pure red noise, is not rejected if

$$\max_{1 \leq k \leq H} |\hat{p}_k - \mu_k| / \tilde{\sigma}_k < q;$$

otherwise, the null hypothesis is rejected and a signal is detected.

5. If H_0 is rejected, then post-hoc testing can be performed for a fixed k : the contribution of W_k (and of the corresponding frequency) is significant if $|\hat{p}_k - \mu_k| / \tilde{\sigma}_k$ exceeds q . Thus, $[\max(0, \mu_k - qw_k \sigma_k), \mu_k + qw_k \sigma_k]$ are considered as the corrected prediction intervals, $k = 1, \dots, H$.

Algorithms 2 and 4 correspond to two-sided prediction intervals (this is the conventional version of Monte Carlo SSA), while Algorithms 1 and 3 correspond to the one-sided case. This yields hyper-rectangular and half hyper-rectangular prediction regions, respectively.

Remark 2. The choice of the vectors W_k , $k = 1, \dots, H$, was discussed in Section 3.4. To increase the criterion power, the number H of the vectors should be as small as possible, e.g., only vectors with dominant frequencies from a given frequency range can be taken. If the vectors W_k are sine waves, the choice of vectors with frequencies from the given range is trivial. If W_k are the eigenvectors, then their dominant frequencies can be calculated by e.g. the ESPRIT method [6, Section 3.1]. This has a little sense for white noise, for which each eigenvector is a mixture of a lot of frequencies. However, for red noise, it has sense, since the eigenvectors correspond to narrow ranges of frequencies. As we mentioned above, increasing the power can be performed by setting the weights.

4. NUMERICAL INVESTIGATION

Let us introduce abbreviations for the test versions. We consider 3-symbols abbreviations with an optional information

$$(3) \quad \{M,S,B\}\{T,E\}\{1,2\}[\text{Est}].$$

The first letter M (Multiple) corresponds to Algorithms 3 and 4, S (Single) corresponds to Algorithms 1 and 2 without Bonferroni correction, while B means using this correction. The second letter is related to the way of generating the projection vectors: T means that the eigenvectors of the theoretical covariance matrix (1) of red noise ($C = 0$) are considered; E means that the eigenvectors of the empirical covariance matrix $\mathbf{X}\mathbf{X}^T$ are used. The last digit means what (one- or two-tailed) test is considered. We consider the method MT1 as the basic one, which is exact and more powerful.

This abbreviation corresponds to the use of true parameters of red noise and unit weights w_k . If the noise parameters are estimated, then we will add ‘Est’ after the digit. If another modification is used, this is indicated separately.

R-scripts in [5] contain an implementation of the above algorithms.

Let us demonstrate the results of different versions of Monte Carlo SSA.

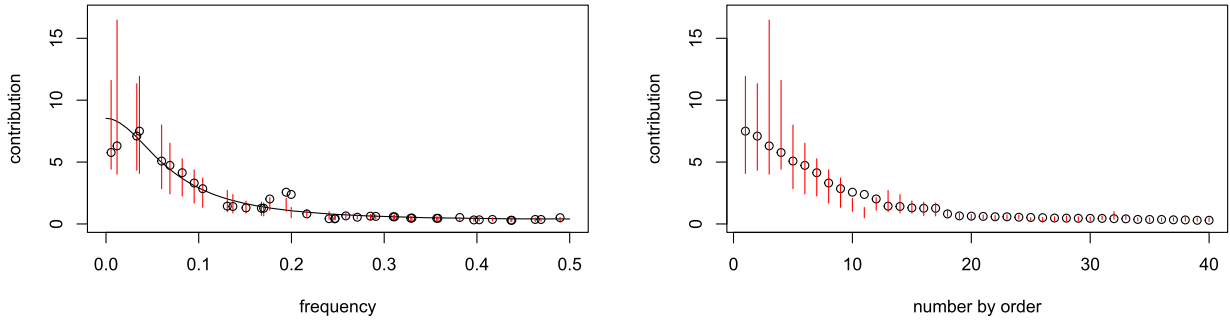


Figure 1. ME2. Left — ordered by frequency, right — ordered by projection value (\hat{p}_k).




Figure 2. $\{S,B\}T2$. Left — without correction, right — with Bonferroni correction.

4.1 Example

The model of an artificial time series is

$$(4) \quad x_n = A \sin(2\pi\omega n) + \xi_n,$$

where ξ_n is red noise with parameters φ and δ , $n = 1, \dots, N$. The case $A = 0$ corresponds to the null hypothesis and the case $A > 0$ yields the presence of signal, that is, corresponds to an alternative. Hereafter we consider the AR(1) parameters $\varphi = 0.7$ and $\delta = 1$.

For illustrative examples, we take $N = 1000$ and the signal parameters $A = 0.4$ and $\omega = 0.2$ in H_1 . The parameters of MC-SSA are $L = 40$ and $G = 1000$.

Let us consider a fixed significance level 0.2 (that is, the confidence level equals 0.8). To weaken the dependence on the time series length N , we consider $\|\mathbf{X}\|^2 = \sum_{i=1}^N x_i^2/N$. We used the true parameters of AR(1) for the creation of surrogate data. The continuous curve is the spectral density of AR(1) with the parameters that were used in the simulation. We calculated the dominant frequencies of W_k by the ESPRIT method with the rank $r = 2$.

Although one-tailed versions of the tests are more powerful, on most of the figures we show two-sided prediction intervals for a clearer visual presentation.

The presence of a sine-wave signal mostly corresponds to exceeding the upper bound. Therefore, as a rule, for exact tests, if the upper bound is smaller, then the test is more powerful.

Fig. 1 demonstrates generally a liberal test, which therefore cannot be used in practice as is. In this version the projection vectors are generated by the observed time series; this choice corresponds to a popular version of Monte Carlo SSA.

Fig. 2 shows prediction intervals for the single test with projection vectors chosen as theoretical eigenvectors; the Bonferroni correction is used in Fig. 2 (right). The test in Fig. 2 (left) is strongly liberal, while Fig. 2 (right) is slightly conservative.

Fig. 3 contains multiple prediction intervals, which provide the exact test with theoretical eigenvectors as the projection vectors. In Fig. 3 (left), the eigenvectors of the matrix (1) with $C = 0$ are taken as W_k . The upper bounds are slightly lower than in Fig. 2 (left). In Fig. 3 (right), the information about the presence of a signal of frequency $\omega = 0.2$ is used and the eigenvectors of the matrix (1) with $C = 1$ are taken as W_k .

Figs. 4 and 5 are related to the case of zero weights for the frequencies of the projection vectors outside the frequency interval $[0.1, 0.3]$. Fig. 4 demonstrates the difference between two-sided and one-sided prediction intervals. In the one-sided case, the upper bound is slightly lower. Fig. 5 shows how the weights influence the prediction intervals sizes. We set $\mathbf{weights} = \mathbf{c}(\text{seq}(6.25, 8, 0.25), 8:1)$, what means that the weights increase from 6.25 to 8 with step 0.25 and then decrease from 8 to 1. Influencing weights on the test power will be studied later.

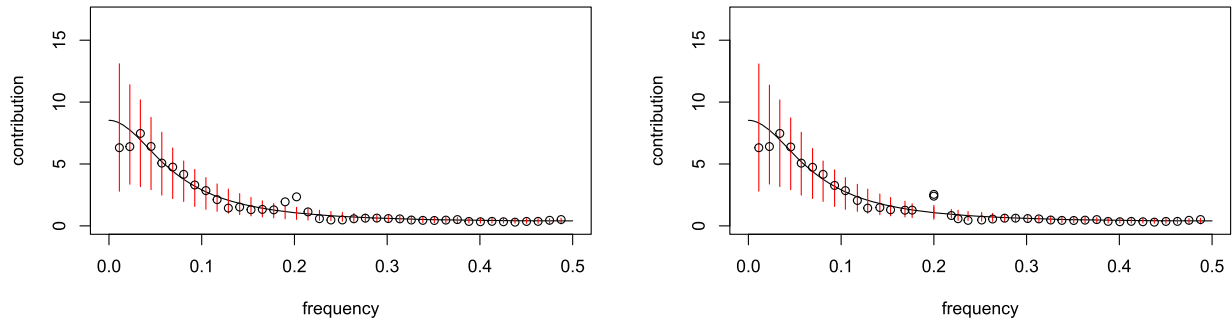


Figure 3. MT2. Left — eigenvectors of pure red-noise covariance matrix, right — eigenvectors of the covariance matrix of red-noise with an added sinusoid at frequency 0.2.




Figure 4. $MT\{1,2\}$, frequency range $[0.1, 0.3]$. Left — two-tailed test, right — one-tailed test.




Figure 5. MT2, weighted, frequency range $[0.1, 0.3]$, two-tailed test. Left — the weights to obtain equal prediction intervals ($w_k = 1/\sigma_k$), right — the weights to increase the power against frequencies around 0.2.

4.2 Study of statistical properties of MC-SSA

Let us describe the methodology of the study of statistical properties of the constructed criteria with bootstrapping. The key items are:

1. The first step is to simulate synthetic data according to the null hypothesis M times and estimate $\alpha_I = \alpha_I(\alpha)$ as the proportion of the rejected null hypothesis for a given significant level α . Check that the necessary condition $\alpha_I \leq \alpha$ is fulfilled. It is essential to choose a sufficient size G of the surrogate data.
2. If the necessary condition is fulfilled, then the second step is to simulate synthetic data according to an alternative hypothesis M times and estimate the power $1 - \alpha_{II}$ against this hypothesis as the proportion of the rejected null hypothesis for a given significant level α .
3. Compare different criteria by power against the alternative hypothesis under interest and use the one with the larger power.
4. If there is no test with sufficient power, then it is possible to improve a test, where $\alpha_I < \alpha$ or $\alpha_I > \alpha$. Then find $\tilde{\alpha}$ such that $\alpha_I(\tilde{\alpha}) \approx \alpha$ and use the significance level $\tilde{\alpha}$ instead of α .




Figure 6. Type I error for different L , empirical and theoretical eigenvectors. Left — the multiple test, right — the single test with the Bonferroni correction.

Below we consider these items. As before, we consider the AR(1) parameters $\varphi = 0.7$ and $\delta = 1$. In a case when the AR(1) parameters are estimated, the maximum likelihood method was applied for estimation, where the conditional-sum-of-squares method was used to find starting values.

Choice of G . A sufficient value G in item 1 can be determined theoretically. The value of G should be considerably large to estimate the quantiles for the surrogate data. In the case of single prediction intervals with Bonferroni correction, $(1 - \alpha/H)$ -quantiles should be estimated (let us consider one-tailed tests), while in the case of multiple testing, we need $(1 - \alpha)$ -quantiles only. For example, if $\alpha = 0.05$, for multiple testing, $G = 1000$ is enough, since the estimate is the 50-th value from the maximum in an ordered sample. However, for single testing and $H = 100$ projection vectors, $\alpha/100 = 0.005$ and therefore the 0.995-quantile will be underestimated (and therefore the test will be liberal). Table 1 contains the estimates of the type I errors for $N = L = 200$ and $\alpha = 0.2$ (then the Bonferroni correction gives the significance level 0.001); $M = 10000$ simulations were used. Table 1 confirms that the multiple methods need smaller G . Hereinafter, the columns ‘2.5%’ and ‘97.5%’ show the bounds of 95%-confidence intervals for the estimated probabilities (the type I error or the power).

Table 1. Type-I error for different G

Method	G	est. type I error	2.5%	97.5%
MT1	100	0.466	0.456	0.476
BT1	100	0.864	0.857	0.870
MT1	500	0.250	0.241	0.258
BT1	500	0.326	0.316	0.335
MT1	1000	0.219	0.211	0.227
BT1	1000	0.212	0.204	0.221

Type I error. The estimates of the family-wise type I errors are contained in Fig. 6; we use $N = 100$, $M = 1000$, and $G = 1000$.

Fig. 6 (left) demonstrates the difference between the use of theoretical and empirical eigenvectors. The pluses and minuses of these versions are as follows.

1. Multiple test, empirical eigenvectors.
 - Plus: It is possible to reconstruct the signal on the basis of significant projection vectors (empirical eigenvectors).
 - Minus: The test is liberal.
2. Multiple test, theoretical eigenvectors.
 - Plus: The test is exact.
 - Minus: If the SSA reconstruction of a detected signal should be performed, a paring is needed to put into correspondence theoretical and empirical eigenvectors.

The comparison of Fig. 6 (right) and Fig. 6 (left) shows the difference between multiple and Bonferroni-corrected single tests. The pluses and minuses of multiple and single versions are as follows.

1. Multiple test
 - Plus: The test is exact.
 - Minus: Recalculation is necessary to consider a subset of the projection vectors.
2. Single test
 - Plus: It is easy to consider subsets of projection vectors without recalculation.
 - Minus: For the version without correction, it is a very liberal test. For the version with Bonferroni correction, it is a slightly conservative test (this is visible for small L) and large G is needed (visible for large L).

In addition, Fig. 6 shows that if the noise parameters are estimated, then the MT and BT tests becomes very conservative (the type I error is considerably smaller than the given level $\alpha = 0.2$). Also, the conservativeness of Bonferroni correction can compensate for the liberality of the test {M,B}E; it can be a reason why the BE test is a bit less liberal than the ME one in Fig. 6.

Power. Let us estimate the power of different versions of the Monte Carlo SSA tests. Consider the presence of a signal in H_1 with $A = 1$ and study the dependence of test power on the signal frequency. We used a standard method of estimation of AR(1) parameters, which ignores a possible presence of a signal.

The estimates of power are depicted in Fig. 7. One can see that both the one-tailed version and a narrow frequency interval can increase the test power.




Figure 7. Power for alternatives with sinusoidal signals of different frequencies.

Correction of type I error. Let us illustrate item 4 of the described scheme with the help of the example (4), where H_0 corresponds to $A = 0$ and H_1 corresponds to $A = 1$, $\omega = 0.1$. As before, we consider a fixed significance level 0.2.

If the noise parameters are estimated, then the criterion becomes very conservative (the type I error is two times smaller than the given $\alpha = 0.2$). Changing α , we can find that for $\tilde{\alpha} = 0.35$ we obtain $\alpha_I \approx 0.2 = \alpha$ (Table 2, the first three rows).

The criterion with the adjusted formal significance level is more powerful than the original conservative criterion. The last three rows of Table 2 show that the use of the conservative test MT1est decreases the power in comparison with that of MT1, while after correction of the significance level to $\tilde{\alpha} = 0.35$, the power of MT1est considerably increases.

Table 2. Improvement of the power of a conservative test

	est. type I error	2.5%	97.5%
MT1	0.200	0.176	0.226
MT1est($\tilde{\alpha} = 0.2$)	0.069	0.054	0.086
MT1est($\tilde{\alpha} = 0.35$)	0.208	0.183	0.234
	est. power	2.5%	97.5%
MT1	0.800	0.774	0.824
MT1est($\tilde{\alpha} = 0.2$)	0.575	0.544	0.606
MT1est($\tilde{\alpha} = 0.35$)	0.731	0.702	0.758

Note that the dependence of the test power on the significance level shows the test strength against a chosen alternative hypothesis. To find how the test would work if the formal significance level is corrected to obtain a given probability of type I error the dependence of the test power on the probability of type I error should be studied; this dependence is a kind of ROC curve.

5. CONCLUSION

In this paper, we studied Monte Carlo SSA from a statistical viewpoint. A scheme for checking the MC-SSA algorithms for correct family-wise type I errors and comparing the algorithms by power was proposed to avoid wrong conclusions in practice. We demonstrated this scheme on a basic version of MC-SSA. However, it can be applied in a general case.

The numerical comparison of multiple and single versions (with Bonferroni correction), of different choices of the projection vectors, and of two-tailed and one-tailed tests was performed. Also, we considered weighted versions of the multiple MC-SSA. Our recommendation is to use the multiple one-tailed MC-SSA; weights are applied if information about the possible frequencies is available. Also, we propose to correct the significance level if the used test is liberal or conservative.

ACKNOWLEDGMENT

We are grateful to Alex Shlemov for the fruitful discussions.

Received 1 March 2021

REFERENCES

- [1] ALLEN, M. R. and ROBERTSON, W. A. (1996). Distinguishing modulated oscillations from coloured noise in multivariate datasets. *Climate Dynamics* **12** 775–784.
- [2] ALLEN, M. R. and SMITH, L. A. (1996). Monte Carlo SSA: Detecting irregular oscillations in the Presence of Colored Noise. *Journal of Climate* **9** 3373–3404.
- [3] BOYAROV, A. (2012). Study of statistical properties of Monte Carlo SSA Master’s thesis, St. Petersburg State University (Scientific Supervisor Nina Golyandina; in Russian).
- [4] GARNOT, V. S. F., GROTH, A. and GHIL, M. (2018). Coupled Climate-Economic Modes in the Sahel’s Interannual Variability. *Ecological Economics* **153** 111–123.
- [5] GOLYANDINA, N. (2021). R-scripts for multiple testing in Monte Carlo singular spectrum analysis. <https://zenodo.org/record/4568495>.
- [6] GOLYANDINA, N., KOROBAYNIKOV, A. and ZHIGLJAVSKY, A. (2018). *Singular spectrum analysis with R*. Springer-Verlag Berlin Heidelberg. [MR3793637](https://doi.org/10.1007/978-3-662-57936-7)
- [7] GOLYANDINA, N., NEKRUTKIN, V. and ZHIGLJAVSKY, A. (2001). *Analysis of Time Series Structure: SSA and Related Techniques*. Chapman & Hall/CRC. [MR1823012](https://doi.org/10.1002/978111823012)
- [8] GRECO, G., ROSA, R., BESKIN, G., KARPOV, S., ROMANO, L., GUARNIERI, A., BARTOLINI, C. and BEDOGNI, R. (2011). Evidence of deterministic components in the apparent randomness of GRBs: clues of a chaotic dynamic. *Scientific reports* **1** 91.
- [9] GRENANDER, U. and SZEGÖ, G. (1984). *Toeplitz Forms and their Applications*. Chelsea, New York. [MR0890515](https://doi.org/10.1002/978111823012)
- [10] JEMWA, G. T. and ALDRICH, C. (2006). Classification of process dynamics with Monte Carlo singular spectrum analysis. *Computers & Chemical Engineering* **30** 816–831.
- [11] KOROBAYNIKOV, A., SHLEMOV, A., USEVICH, K. and GOLYANDINA, N. (2021). Rssa: A collection of methods for singular spectrum analysis R package version 1.0.4.
- [12] LANCASTER, G., IATSENKO, D., PIDDE, A., TICCINELLI, V. and STEFANOVSKA, A. (2018). Surrogate data for hypothesis testing of physical systems. *Physics Reports* **748** 1–60. [MR3830658](https://doi.org/10.1016/j.physrep.2018.08.001)

- [13] LOUCA, S. and DOEBELI, M. (2015). Detecting cyclicity in ecological time series. *Ecology* **96** 1724–1732.
- [14] PALUS, M. and NOVOTNÁ, D. (1998). Detecting modes with non-trivial dynamics embedded in colored noise: enhanced Monte Carlo SSA and the case of climate oscillations. *Physics Letters A* **248** 191–202.
- [15] PALUS, M. and NOVOTNÁ, D. (2004). Enhanced Monte Carlo Singular System Analysis and detection of period 7.8 years oscillatory modes in the monthly NAO index and temperature records. *Nonlinear Processes in Geophysics* **11** 721–729.
- [16] TRENCH, W. F. (2010). Spectral decomposition of Kac–Murdoch–Szegő matrices.
- [17] VAUGHAN, S. (2005). A simple test for periodic signals in red noise. *A&A* **431** 391–403.

Nina Golyandina
Department of Statistical Modeling
St. Petersburg State University
Universitetskaya nab. 7/9
St. Petersburg
Russia
E-mail address: n.golyandina@spbu.ru