Journal of Computational and Applied Mathematics 436 (2024) 115379

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Stationary subspace analysis based on second-order statistics

Lea Flumian [a], Markus Matilainen [b,c], Klaus Nordhausen [d,*], Sara Taskinen [d]

[a] 22 Boulevard de Reuilly, 75012 Paris, France
[b] Turku PET Centre, Turku University Hospital and University of Turku, Finland
[c] Faculty of Science and Engineering, Åbo Akademi University, Turku, Finland
[d] Department of Mathematics and Statistics, University of Jyväskylä, Finland

ARTICLE INFO

Article history:
Received 24 January 2022
Received in revised form 4 May 2023

Keywords:
Autocorrelation
Joint diagonalization
Multivariate time series
Second-order stationary
Supervised dimension reduction

ABSTRACT

In stationary subspace analysis (SSA) one assumes that the observable p-variate time series is a linear mixture of a k-variate nonstationary time series and a $(p-k)$-variate stationary time series. The aim is then to estimate the unmixing matrix which transforms the observed multivariate time series onto stationary and nonstationary components. In the classical approach multivariate data are projected onto stationary and nonstationary subspaces by minimizing a Kullback–Leibler divergence between Gaussian distributions, and the method only detects nonstationarities in the first two moments. In this paper we consider SSA in a more general multivariate time series setting and propose SSA methods which are able to detect nonstationarities in mean, variance and autocorrelation, or in all of them. Simulation studies illustrate the performances of proposed methods, and it is shown that especially the method that detects all three types of nonstationarities performs well in various time series settings. The paper is concluded with an illustrative example.

1. Introduction

Multivariate time series data are observed in many application areas and are often very challenging to model. To ease the analyses, it is often assumed that the observed time series can be decomposed into latent components with different exploitable properties. One common approach is to apply one of the blind source separation (BSS) methods to observed data in order to estimate the latent components as a pre-processing tool. In a special case of BSS, that is, in second-order source separation (SOS), it is assumed that the latent components are uncorrelated second-order stationary time series, and in independent component time series model, the assumption on uncorrelatedness is replaced with the assumption on independence (see *e.g.* [1], and references therein). In nonstationary source separation (NSS), as defined in [2], the variances of uncorrelated sources are allowed to change over time while the mean is assumed to be stationary. For recent reviews of these and other variants of BSS, see [3–7].

In stationary subspace analysis (SSA), the underlying model states that the observable p-variate time series is a linear mixture of a k-variate nonstationary time series and a $(p-k)$-variate stationary time series. The aim is then to factorize the multivariate time series onto stationary and nonstationary components. Such a factorization is useful in several real world applications, *e.g.*, in biomedical signal processing, speech recognition, image analysis and econometrics, where either the stationary or the nonstationary components are of interest. The method development was motivated by the EEG data analysis and computer vision applications [8–10]. Later, [11,12] applied SSA to change point detection which one of the

* Corresponding author.
E-mail address: klaus.k.nordhausen@jyu.fi (K. Nordhausen).

https://doi.org/10.1016/j.cam.2023.115379

fields where nonstationary components are also of the interest. Recall that in the detection of change points, also known as structural changes, one searches for time points where the dynamics of the underlying stochastic process changes. Change point detection is an active research area for example in the financial time series analysis (for a review, see for example Aue and Horvath [13]), and it can benefit from the SSA methodology as the problem becomes much easier to solve in a lower-dimensional space as there is no need to search for change points in the stationary subspace. For other examples of financial time series that are nonstationary in mean, variance and in autocorrelation, we refer to [14,15], and references therein. Finally notice that another family of methods that aims at finding stationary components is based on costationarity [16]. The method however differs from the classical SSA approach in that is assumes that stationary components are time-varying linear combinations of locally stationary time series.

SSA was introduced in [8], where the matrix that separates the subspaces of stationary and nonstationary components was found by dividing the observed multivariate time series into K segments and minimizing a Kullback–Leibler (KL) divergence between Gaussian distributions thus measuring differences in means and covariances across segments. We denote such method as KL-SSA. Notice that in KL-SSA the time series are considered as stationary if the first two moments are time-invariant. The stationarity with respect to the autocorrelations is not taken into account. [8] discussed the theoretical properties of the proposed method and proposed a sequential likelihood ratio test for testing the dimension of the stationary subspace. In [17,18] several alternatives for the Kullback–Leibler divergence in SSA were given.

In this paper we propose a novel SSA algorithm that detects nonstationarities in the first two moments as well as in autocorrelations. The method can be seen as an extension of analytic SSA (ASSA), which is a SSA method suggested in [19], to a more general multivariate time series setting. Unlike ASSA, our proposed method provides a tool to identify the type of nonstationarities it detects. As another contribution in this paper we make a connection between SSA and supervised dimension reduction (SDR) methods. Notice that [20] proposed an alternative extension to ASSA that aims to detect nonstationarities in mean, variance and autocorrelations using a frequency domain approach. We refer to their method as dependent-SSA (DSSA) and compare its performance to the methods proposed in this paper. In [20] a sequential testing procedure for the dimension of the stationary subspace is also provided.

The paper is organized as follows. In Section 2 we first review the model assumptions underlying SSA and discuss the two classical methods for performing SSA, that is, KL-SSA and ASSA. We also show how SSA can be seen as a supervised dimension reduction method and propose and compare methods for detecting different type of nonstationarities. Section 3 discusses practical issues and identifiability of stationary and nonstationary components. Section 4 presents several simulation studies for comparing performances of the proposed methods to performances of ASSA and DSSA under various scenarios. A real data example is also provided. The paper is concluded with some discussion in Section 5.

2. Stationary subspace analysis

In this section we review the model underlying SSA, show some connections between SSA and supervised dimension reductions methods, and suggest methods to detect changes in mean, variance, correlation structure or in all of them. Tools to identify the type of nonstationarity the components exhibit are also given.

2.1. Model formulation and notations

Let $\mathbf{x_t}$ be an observable p-variate nonstationary time series which can be decomposed into a stationary part and a nonstationary part according to the SSA model as defined in [8]

$$\mathbf{x}_t = \mathbf{A}\mathbf{z}_t = [\mathbf{A_s}\ \mathbf{A_n}] \begin{pmatrix} \mathbf{s}_t \\ \mathbf{n}_t \end{pmatrix}, \tag{1}$$

where a p-variate latent time series $\mathbf{z}_t$ consists of a k-variate nonstationary time series $\mathbf{n}_t$ and a $(p-k)$-variate stationary time series $\mathbf{s}_t$. The two components are mixed using a full-rank $p \times p$ matrix $\mathbf{A}$. Here matrices $\mathbf{A_s}$ and $\mathbf{A_n}$ are $p \times (p-k)$ and $p \times k$ matrices, respectively. The aim of SSA is to estimate the unmixing matrix $\mathbf{W} = \mathbf{A}^{-1}$ so that $\mathbf{W}\mathbf{x}_t$ is partitioned into stationary and nonstationary time series. Depending on the assumptions on $\mathbf{s}_t$ and $\mathbf{n}_t$, different ways to estimate $\mathbf{W}$ are suggested in the literature [8,19]. In this paper we make the following assumptions on the latent time series:

(A1) $E(\mathbf{s}_t) = \mathbf{0}$, $\mathrm{Cov}(\mathbf{s}_t) = \mathbf{I}_{p-k}$ and $\mathrm{Cov}(\mathbf{s}_1, \mathbf{s}_{1+\tau}) = \mathrm{Cov}(\mathbf{s}_t, \mathbf{s}_{t+\tau}) < \infty$,
(A2) $E(\mathbf{n}_t) < \infty$, $\mathrm{Cov}(\mathbf{n}_t) = \mathbf{D}_t$, where $\mathbf{D}_t$ is a diagonal matrix with positive diagonal elements, and $\mathrm{Cov}(\mathbf{n}_t, \mathbf{n}_{t+\tau}) < \infty$,
(A3) The latent time series $\mathbf{s}_t$ and $\mathbf{n}_t$ are independent, i.e., $\mathrm{Cov}(\mathbf{s}_t, \mathbf{n}_{t+\tau}) = \mathbf{0}$ for all $\tau \geq 0$.

We thus assume that the $(p-k)$ time series in $\mathbf{s}_t$ are second-order stationary with finite autocovariances. The first assumption fixes the location and covariance matrix of the stationary part for convenience. The second assumption states that the k nonstationary components in $\mathbf{n}_t$ have finite first and second moments and are uncorrelated at lag $\tau = 0$ but can have arbitrary time dependent autocovariances at all other lags. The third assumption states that nonstationary and stationary components are independent.

Despite the assumptions the model is not well-defined as there are many divisions of $\mathbf{x}_t$ such that the assumptions hold. Here we are interested in the one with the minimal value of k. The stationary components are only specified up to a rotation by an orthogonal matrix, whereas the nonstationary components can be marginally rescaled, shifted and also rotated. Therefore for convenience of presentation we make the additional assumption on $\mathbf{n}_t$:

(A4) $\sum_{t\in T} E(\mathbf{n}_t) = \mathbf{0}$ and $\sum_{t\in T} \mathrm{Cov}(\mathbf{n}_t) = \sum_{t\in T} \mathbf{D}_t = \mathbf{I}_k$.

Hence for the observed time span T, location and scale are considered fixed.

The common idea in [8,19] was to divide the observed time series into K segments, then compute for each interval the first and the second order statistics and to measure their nonconstancy across segments. To be more exact, assume that $\mathbf{x}_t$ is observed at time points $1, \ldots, T$, and let $T_1, \ldots, T_K$ denote K disjoint subsets of T, *e.g.*, non-overlapping intervals. Let then

$$\mathbf{m}_{T_i}(\mathbf{x}_t) = \frac{1}{|T_i|}\sum_{t\in T_i} \mathbf{x}_t$$

and

$$\mathbf{S}_{\tau,T_i}(\mathbf{x}_t) = \frac{1}{|T_i|-\tau}\sum_{t\in T_i} (\mathbf{x}_t - \mathbf{m}_{T_i}(\mathbf{x}_t))(\mathbf{x}_{t+\tau} - \mathbf{m}_{T_i}(\mathbf{x}_t))^\top$$

denote the sample mean and the sample (auto)covariance matrix computed using the data from interval T_i, respectively. In [8,19] the unmixing matrix revealing the stationary sources was chosen so that the means and covariance matrices vary the least across all intervals. It was also shown that using the standardized time series defined as

$$\mathbf{y}_t = \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}(\mathbf{x}_t - \mathbf{m}_T(\mathbf{x}_t)), \tag{2}$$

the search can be restricted to orthogonal matrices only. In [8] KL divergence between the Gaussian distributions on each intervals and the standard normal distribution was minimized. Further, in [19] a second-order Taylor approximation was applied to the objective function of KL-SSA and the SSA problem was shown to reduce to a simple generalized eigenvalue problem. We review the resulting analytic SSA (ASSA) method in Section 2.7. Notice that in both approaches the focus was on detecting stationarity deviations in mean and in variance but not in autocorrelation.

2.2. Connection to supervised dimension reduction

Before going into specific SSA methods we point out a connection between SSA and supervised dimension reduction (SDR). In SDR, we have a response z and a p-variate predictor vector $\mathbf{x}$ and the goal is to find a $k \times p$ matrix $\mathbf{W}$ such that $\mathbf{W}^\top\mathbf{x}$ carries all relevant information about z, that is, $\mathbf{x} \perp\!\!\!\perp z \mid \mathbf{W}^\top\mathbf{x}$ with the smallest possible value of $k \ll p$, see [21,22] for some general overviews. Two popular methods in this context are sliced inverse regression [SIR, 23] and sliced average variance estimation [SAVE, 24]. For a sample $(z_j, \mathbf{x}_j)$, $j = 1, \ldots, n$, both methods start by whitening $\mathbf{x}_j$, that is, by computing $\mathbf{y}_j = \mathbf{S}_{0,n}^{-1/2}(\mathbf{x}_j - \mathbf{m}_n(\mathbf{x}_j))$ and then group the whitened observations into K so-called slices $H_1, \ldots, H_K$ according to their response values z_i. By defining the slice means as $\mathbf{m}_{H_i}(\mathbf{y}) = |H_i|^{-1}\sum_{i\in H_i}\mathbf{y}_i$ and slice covariance matrices as $\mathbf{S}_{H_i}(\mathbf{y}) = |H_i|^{-1}\sum_{i\in H_i}(\mathbf{y}_i - \mathbf{m}_{H_i}(\mathbf{y}))(\mathbf{y}_i - \mathbf{m}_{H_i}(\mathbf{y}))^\top$, where $i = 1, \ldots, K$, SIR estimate is then obtained using the covariance of the slice means, $n^{-1}\sum_{i=1}^K |H_i|\,\mathbf{m}_{H_i}(\mathbf{y})\mathbf{m}_{H_i}^\top(\mathbf{y})$, and SAVE estimate is based on the average of the differences between the slice covariance matrices and the covariance matrix of all $\mathbf{y}_i$, that is, $n^{-1}\sum_{i=1}^K |H_i|\,(\mathbf{I}_p - \mathbf{S}_{H_i}(\mathbf{y}))^2$.

If the response variable z_i is categorical, slicing is usually done by the unique values of z_i. In the case of numerical z_i, the slices are often chosen so that they contain approximately the same number of observations. Choosing the number of slices is a trade-off between trying to have many slices to find the directions of interest and to have enough data points in slices in order to estimate the slice statistics sufficiently well. [25] show that while SAVE is more comprehensive than SIR in detecting the subspace of interest, the increased flexibility comes at the cost of requiring more data and being more sensitive to the number of slices. [25] also point out situations where SIR fails. Thus, while SAVE is theoretically superior, the SIR method is in practice often preferred as it is more robust with regard to the selection of number of slices. Also combinations of SIR and SAVE are regularly investigated, see for example [26–28] for more details.

In SDR as described above there is a natural response variable z which is not available in the context of SSA. In SSA we basically construct an artificial response by dividing the time series into intervals and the aim is then to find the components which carry information about interval "membership". Then, as the stationary components do not change their behavior over time, they carry no information regarding the intervals. The nonstationary components on the other hand change over time and therefore provide information that help in identifying the response. In the following sections we consider SSA methods that correspond to sliced inverse regression [SIR, 23] and sliced average variance estimation [SAVE, 24], respectively. Section 2.5 proposes a method for detecting stationarity deviations in autocorrelation as these cannot be detected by SIR and SAVE type approaches. Further, inspired by the success of hybrid methods in SDR we also consider a combined approach in Section 2.7.

Note that SDR methods, such as SIR and SAVE, and combinations of these have also been extended to the time series settings, see [29,30] who call the methods accordingly TSIR and TSAVE. These methods however focus on regression modeling and not for detecting nonstationary subspaces. Nevertheless, especially in the time series context, TSIR and TSAVE have been considered also in a latent factor model similar to ours, but the same type of latent factor assumptions

can also be made for SIR and SAVE [31]. Meaning they derived the methods in a framework where the explaining variables $\mathbf{x}$ can also be presented as

$$\mathbf{x} = \mathbf{A}\mathbf{z},$$

where $\mathbf{A}$ is a full rank matrix $\mathbf{z} = (\mathbf{z}_i^\top, \mathbf{z}_2^\top)^\top$ is a random p-vector with $\mathbb{E}(\mathbf{z}) = \mathbf{0}$, $\mathrm{Cov}(\mathbf{z}) = \mathbf{I}_p$. Then the key requirement in this paper for SIR and SAVE are that $(z, \mathbf{z}_1^\top)^\top \perp\!\!\!\perp \mathbf{z}_2$, meaning for example that $\mathbf{z}_2$ carries no information about the response z. This means that at the population

$$\mathrm{Cov}(E(\mathbf{z}|y)) = \begin{pmatrix} \mathrm{Cov}(E(\mathbf{z}_1|y)) & 0 \\ 0 & 0 \end{pmatrix}$$

and

$$E(\mathbf{I}_p - \mathrm{Cov}(\mathbf{z}|y)) = \begin{pmatrix} E(\mathbf{I}_k - \mathrm{Cov}(\mathbf{z}_1|y)) & 0 \\ 0 & 0 \end{pmatrix}$$

have block diagonal form, which together with the affine equivariance of the methods are the key ingredients of SIR and SAVE. Note that our requirement stated here is stronger than the original conditions imposed on SIR (linearity condition) and SAVE (constant variance condition). For a discussion about these assumptions we refer to [31] and references therein. Our assumptions here seem however more natural in a time series context as argued in Matilainen et al. [29,30] and translate naturally into the SSA framework. This means that if we denote z_t as the variable that indicates the interval membership, we have accordingly $(z_t, \mathbf{n}_t^\top)^\top \perp\!\!\!\perp \mathbf{s}_t$ as an assumption. This will guarantee that under proper selection of the intervals the key matrices of interest, which will be defined in the following, have then also the same block diagonal form.

2.3. Nonstationarity in mean

Consider first a SSA method that aims at detecting stationary deviations in mean. Examples of such nonstationary components include trend, seasonal and cyclic components. Following [8,19] we choose the unmixing matrix $\mathbf{W}$ revealing the stationary (nonstationary) components so that the means vary the least (most) across all intervals. Now the nonconstancy in mean is measured using SIR type covariance matrix of means computed on different intervals. Similar to the most BSS methods, the use of the standardized, whitened time series (2) when computing the covariance matrix guarantees that the search can be restricted to orthogonal matrices only [4].

Assume now that $\mathbf{x}_t$ is generated from the SSA model (1) and write $\mathbf{y}_t$ for the standardized time series. Let then

$$\mathbf{M}_m = \sum_{i=1}^{K} \frac{|T_i|}{T} \mathbf{m}_{T_i}(\mathbf{y}_t)\mathbf{m}_{T_i}(\mathbf{y}_t)^\top \tag{3}$$

be the covariance matrix of means computed on different intervals weighted by the interval length. Notice that with seasonal and cyclic components, the intervals have to be chosen so that they do not correspond to the cycles. This is also illustrated in Section 2.6. If the mean of a univariate time series does not change in time, corresponding diagonal element of $\mathbf{M}_m$ is zero. Thus, those components, that are stationary in mean, correspond to zero eigenvalues of $\mathbf{M}_m$. Further, the components that are nonstationary in mean correspond to non-zero eigenvalues. Hence write the eigenvalue decomposition of $\mathbf{M}_m$ as

$$\mathbf{M}_m = \mathbf{U}_m \mathbf{D}_m \mathbf{U}_m^\top,$$

where $\mathbf{D}_m$ is a $p \times p$ diagonal matrix with the eigenvalues of $\mathbf{M}_m$ as diagonal values, and $\mathbf{U}_m = (\mathbf{U}_{m,s}\ \mathbf{U}_{m,n})$ includes the eigenvectors of $\mathbf{M}_m$ arranged so that $\mathbf{U}_{m,n}$ is the $p \times k_m$ matrix containing the eigenvectors belonging to the non-zero eigenvalues as columns and $p \times (p - k_m)$ matrix $\mathbf{U}_{m,s}$ includes the remaining ones as columns. The columns of resulting unmixing matrices $\mathbf{W}_{m,n} = \mathbf{U}_{m,n}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$ and $\mathbf{W}_{m,s} = \mathbf{U}_{m,s}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$ then generate the nonstationary and stationary subspaces, respectively. Naturally $k_m \leq k$ with equality only if for all nonstationary components the means differ at least between some of the chosen intervals. As this method is based on interval means it corresponds basically to SIR and we denote this method accordingly SSAsir.

2.4. Nonstationarity in variance

Following the principles of Section 2.3 we can then derive a SAVE type method for detecting stationary deviations in variance. Consider again standardized time series $\mathbf{y}_t$. Now

$$\mathbf{M}_v = \sum_{i=1}^{K} \frac{|T_i|}{T} (\mathbf{I}_p - \mathbf{S}_{0,T_i}(\mathbf{y}_t))^2, \tag{4}$$

where $\mathbf{A}^2 = \mathbf{A}\mathbf{A}^\top$, measures the deviation of the covariance matrix computed on intervals $T_1, \ldots, T_K$ from the global covariance matrix $\mathbf{I}_p$.

Again, those eigenvalues of $\mathbf{M}_v$ that correspond to the components that are nonstationary in variance should be non-zero. Let thus the eigenvalue decomposition be

$$\mathbf{M}_v = \mathbf{U}_v \mathbf{D}_v \mathbf{U}_v^\top,$$

where $\mathbf{D}_v$ is a $p \times p$ diagonal matrix with the eigenvalues of $\mathbf{M}_v$ as diagonal values, $\mathbf{U}_v = (\mathbf{U}_{v,s}\ \mathbf{U}_{v,n})^\top$ are the eigenvectors of $\mathbf{M}_v$ arranged so that $\mathbf{U}_{v,n}$ is the $p \times k_v$ matrix containing the eigenvectors belonging to the non-zero eigenvalues as columns and $\mathbf{U}_{v,s}$ is the $p \times (p - k_v)$ matrix containing the remaining eigenvectors as columns. Again $k_v \leq k$. The transforming matrices to the two subspaces are accordingly $\mathbf{W}_{v,n} = \mathbf{U}_{v,n}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$ and $\mathbf{W}_{v,s} = \mathbf{U}_{v,s}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$. As illustrated in Section 2.6, although this approach is designed to detect components with nonstationary variances, it may also detect components which are nonstationary in mean as if the mean changes the variances in intervals might also change. Similar phenomenon is seen also in classical SAVE method [25]. We will refer to this method in the following as SSAsave.

2.5. Nonstationarity in autocorrelation

The SIR and SAVE type methods are able to detect components which are nonstationary with respect to the first and the second moments. To detect nonstationarities in autocorrelation we need a statistic to measure the nonconstancy in autocorrelation across segments. Let such statistic be defined (for standardized time series) as

$$\mathbf{M}_\tau = \sum_{i=1}^{K} \frac{|T_i|}{T} (\mathbf{S}_{\tau,T}(\mathbf{y}_t) - \mathbf{S}_{\tau,T_i}(\mathbf{y}_t))^2, \tag{5}$$

where $\tau \geq 1$. The matrix $\mathbf{M}_\tau$ thus measures the deviation of the autocovariance matrix computed on intervals $T_1, \ldots, T_K$ from the global autocovariance matrix $\mathbf{S}_{\tau,T}$. Using again the eigenvalue–eigenvector decomposition

$$\mathbf{M}_\tau = \mathbf{U}_\tau \mathbf{D}_\tau \mathbf{U}_\tau^\top, \tag{6}$$

and separating the eigenvectors of $\mathbf{M}_\tau$ corresponding to non-zero and zero eigenvalues yields the $p \times k_\tau$ matrix $\mathbf{U}_{\tau,n}$ and the $p \times (p - k_\tau)$ matrix $\mathbf{U}_{\tau,s}$, where k_τ is the number of non-zero eigenvalues with $k_\tau \leq k$. The unmixing matrix estimates are $\mathbf{W}_{\tau,n} = \mathbf{U}_{\tau,n}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$ and $\mathbf{W}_{\tau,s} = \mathbf{U}_{\tau,s}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$. For this method to work, the different intervals are required to have different autocovariances in the intervals. As this method is aimed in detecting changes in the correlation structure we denote it as SSAcor. Notice that as $\mathbf{M}_\tau$ is computed using a single lag τ, the performance of SSAcor will naturally depend on the choice of τ. There is however no rule how τ should be selected in practice without having any additional information on the underlying processes. In the BSS literature it is often recommended to start with the choice $\tau = 1$ and then modify the choice if the method shows poor performance [3]. Another option here is to use several lags $\tau_1, \ldots, \tau_L$ and then perform joint diagonalization of $\mathbf{M}_{\tau_1}, \ldots, \mathbf{M}_{\tau_L}$ instead of using the eigenvalue–eigenvector decomposition in (6). We return to this approach in Section 2.7.

2.6. Comparison of SSAsir, SSAsave and SSAcor

In this section we visualize how SSAsir, SSAsave and SSAcor work and in which situations they fail. For that purpose we plot four nonstationary time series of length 3000 which are all standardized to have mean zero and unit variance. The settings of the time series are given in Appendix A. We consider $K = 6$ equal-sized intervals and compute for each slice the mean, the variance and the autocovariance, and visualize these quantities. The three quantities are illustrated in Figs. 1 and 2 using dots, vertical bars and horizontal bars, respectively. The methods now detect nonstationarities when there exists variation between the means (SSAsir), when the variances differ from 1 (SSAsave) which means that the intervals must have different variances, and when the autocovariance differs from the global autocovariance which means that they also need to be different. By combining the values computed at each slice, we then obtain the global statistics of interest for each time series and for each method. These statistics are reported in the top-left corners of the figures and for the methods to work the values need to be non-zero. As can be seen here the values of non-stationary in mean or variance take larger values than those for detecting changes in auto-correlation and are therefore easier to detect. This will also be clear in the simulation studies which will be conducted later. One explanation for this is that the criteria statistics for mean and variance have no upper bounds. However, as the time series here are whitened the autocovariance actually correspond to autocorrelation and therefore has upper bounds.

Fig. 1 shows two different types of nonstationarities in mean. The left panel has a mean level shift and the right panel has a strong seasonal pattern. For the mean level shift the mean values are clearly non-constant but also the variances in the intervals differ. Thus in this case both SSAsir and SSAsave work though the statistic of SSAsir is much larger than that of SSAsave. Notice that SSAcor does not work at all in this case. In case of the seasonal time series the example shows that if the intervals are chosen badly, as is the case here, none of the methods work despite the clear nonstationarity. In Appendix A, Fig. A.13 shows the same examples as in this section with $K = 6$ intervals with unequal sizes. This demonstrates that SSAsir and SSAsave can perform well in the case of seasonal time series and the performance depends on the chosen intervals. Thus, the number of K of intervals as well as their widths are tuning parameters which can influence the performance of the methods.

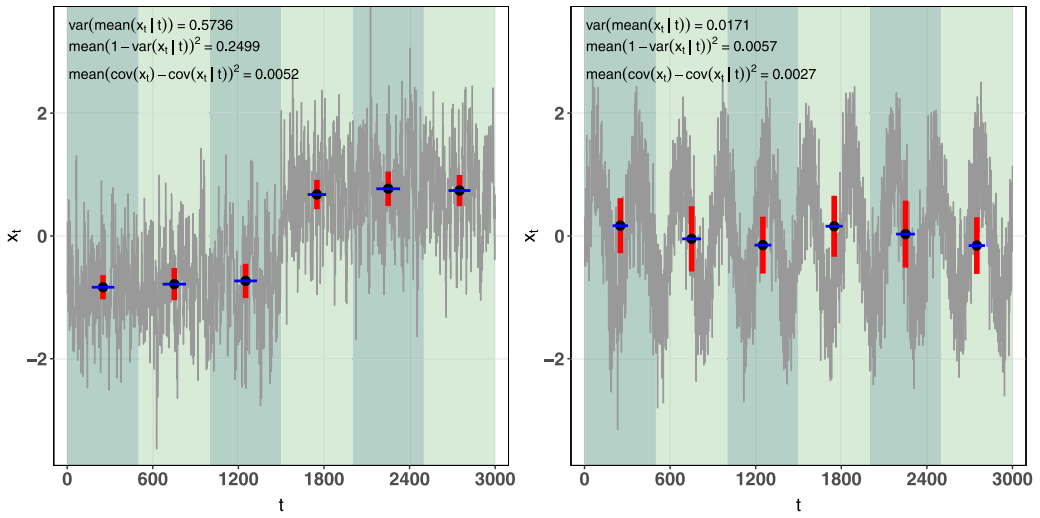

Fig. 1. Visualization of SSAsir, SSAsave and SSAcor for a component with nonstationarity in mean (left panel trend, right panel seasonality). The black dots represent the interval mean, the height of vertical red bar the interval variance and the width of blue horizontal bar the interval autocovariance.

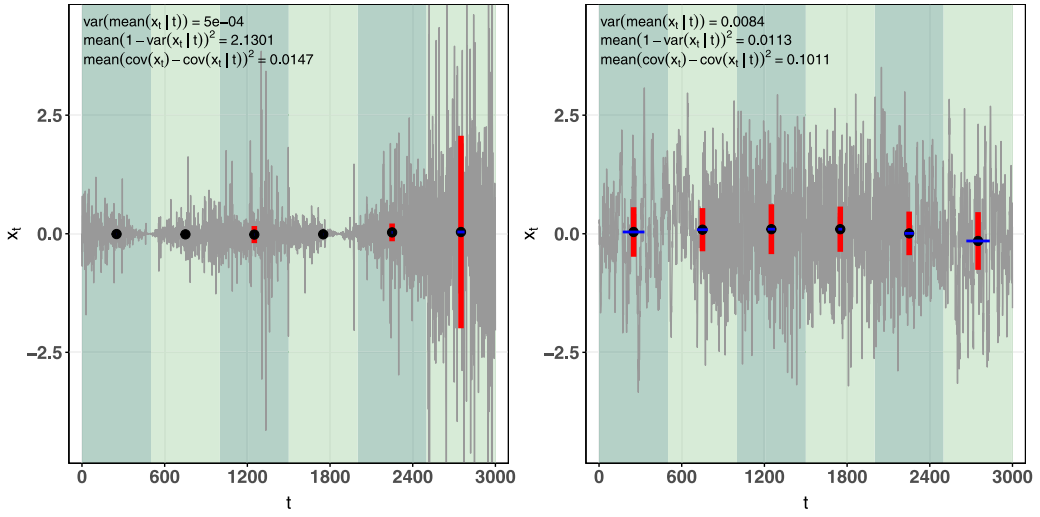

Fig. 2. Visualization of SSAsir, SSAsave and SSAcor for a component with nonstationarity variance (left panel) and for a component with nonstationary correlation structure (right panel). The black dots represent the interval mean, the height of vertical red bar the interval variance and the width of blue horizontal bar the interval autocovariance.

In the left panel of Fig. 2 the time series has a constant mean but changes in variance. As expected SSAsir cannot detect this type of nonstationarity as only varying variances in the intervals are visible. Again SSAcor does not work here. The right panel of Fig. 2 on the other hand has a constant mean and a constant variance but the correlation structure changes twice. This is also visible in the figure where for the first time the horizontal correlation bars appear to have clearly different lengths. However, the overall measure of nonstationarity used by SSAcor is still small indicating that detection of nonstationarity in correlation is quite difficult. As shown in Fig. A.14 in the Appendix having intervals of unequal sizes does not give higher measure on nonstationarity for SSAcor. To conclude, there always seems to be cases where one of the methods is better than the other making approaches that combine the three methods a natural approach.

2.7. Combination of methods

The methods described above can be combined in order to detect all three types of nonstationarities. Notice first that analytic SSA (ASSA) as proposed in [19] tries to recover the unmixing matrix **W** so that the covariance matrices computed

using the first and the second moments across intervals are minimized. To be more exact, the method combines the search for components that are nonstationary in mean and variance by using the eigenvalue–eigenvector decomposition of

$$\mathbf{M}_{ASSA} = \frac{1}{T}\sum_{i=1}^{K}\left\{\mathbf{m}_{T_i}(\mathbf{y}_t)\mathbf{m}_{T_i}(\mathbf{y}_t)^\top + \frac{1}{2}\mathbf{S}_{0,T_i}(\mathbf{y}_t)\mathbf{S}_{0,T_i}(\mathbf{y}_t)^T\right\} - \frac{1}{2}\mathbf{I}_p,$$

and proceeding as in the previous sections. The main drawback here is that the changes in autocorrelation are not necessarily detected. Furthermore there are no tools to identify what kind of nonstationarities the components exhibit.

Therefore, we suggest another approach which combines the three nonstationarity measures (3), (4) and (5). Denote from now on $\mathbf{M}_m = \mathbf{M}_1$, $\mathbf{M}_v = \mathbf{M}_2$ and $\mathbf{M}_\tau = \mathbf{M}_3$. Our suggestion is to turn the problem into a joint diagonalization problem. Notice that the joint diagonalization is a widely used tool for example in second-order source separation [32], where several autocovariance matrices are jointly diagonalized in order to separate stationary sources from their linear mixture. As simultaneous diagonalization is not possible, one has to perform approximate joint diagonalization [33]. Applying this to the three matrices $\mathbf{M}_i$, $i = 1, 2$ and 3, computed using standardized time series means that we search for an orthogonal $p \times p$ matrix $\mathbf{U}_c$ which minimizes $\sum_{i=1}^{3} \|\text{off}(\mathbf{U}_c^\top \mathbf{M}_i \mathbf{U}_c)\|^2$, or, as $\mathbf{U}_c$ is orthogonal, maximizes

$$\sum_{i=1}^{3} \|\text{diag}(\mathbf{U}_c^\top \mathbf{M}_i \mathbf{U}_c)\|^2 \tag{7}$$

[34]. Here $\|\mathbf{A}\|$ is the matrix (Frobenius) norm, $\text{diag}(\mathbf{A})$ is a $p \times p$ diagonal matrix with the diagonal elements as in $\mathbf{A}$ and $\text{off}(\mathbf{A}) = \mathbf{A} - \text{diag}(\mathbf{A})$. Notice that naturally only a subset of the matrices or some additional matrices $\mathbf{M}_\tau$ with various lags τ can be added in the objective function, but the principle of the estimation procedure remains the same. Several algorithms for approximate joint diagonalization in (7) exist in the literature. The most popular one based on Givens rotations is proposed in [35] and is available for example in R package JADE [36].

Based on $\mathbf{U}_c$ one can compute $\mathbf{D}_i = \text{diag}(\mathbf{U}_c \mathbf{M}_i \mathbf{U}_c^\top) = \text{diag}(d_{i,1}, \ldots, d_{i,p})$ and collect all those columns of $\mathbf{U}_c$, where $\sum_i d_{ij} \neq 0$, to a $p \times k_c$ matrix $\mathbf{U}_{c,n}$. The rest of the columns are then collected to a $p \times (p - k_c)$ matrix $\mathbf{U}_{c,s}$. Here the individual value $d_{i,j}$ indicates whether the jth component is nonstationary with respect to $\mathbf{M}_i$. Such classification of the components is not possible when, for example, methods such as ASSA are applied. The final transformation matrices for nonstationary and stationary components are then $\mathbf{W}_{c,n} = \mathbf{U}_{c,n}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$ and $\mathbf{W}_{c,s} = \mathbf{U}_{c,s}^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$. In the following this method is referred as SSAcomb.

3. Practical issues and identifiability of components

There are several practical considerations still worth pointing out. Clearly the finite sample eigenvalues of matrices $\mathbf{M}$ corresponding to the stationary components will be never exactly zero. Thus the value of k must be chosen based on a cut-off value or graphically. Furthermore, it is obvious that the choice of the number of intervals K and how they are divided is crucial in this framework. The impact of the number of intervals will be illustrated in the simulation studies in Section 4. Ideally, given the number of intervals, the cut points of the intervals should be such that the quantities of interest are as different as possible. However, as finding optimal cut points is in practice difficult, the intervals should, in our opinion, be at least of different length so that possible seasonality effects are easier to detect and issues similar to those seen in Section 2.6 do not occur. Thus, the decision regarding the interval length and number of intervals should be based on the data at hand and might be best chosen using visual analytic tools as recently suggested for example in an NSS context in Piccolotto et al. [37].

For all methods discussed above, the unmixing matrix $\mathbf{W}$ has the two parts $\mathbf{W}_n = \mathbf{U}_n^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$ and $\mathbf{W}_s = \mathbf{U}_s^\top \mathbf{S}_{0,T}(\mathbf{x}_t)^{-1/2}$ specifying nonstationary and stationary subspaces, respectively, where matrices $\mathbf{U}_n$ and $\mathbf{U}_s$ depend on the method used. Let now $\mathbf{x}_t^* = \mathbf{B}\mathbf{x}_t$ denote an affine transformation of $\mathbf{x}_t$ with full-rank $p \times p$ matrix $\mathbf{B}$. We then denote the unmixing matrix based on $\mathbf{x}_t^*$ as $\mathbf{W}^*$. In the BSS literature (see for example [38]) a BSS unmixing matrix is defined to be affine equivariant if $\mathbf{W}\mathbf{x}_t = \mathbf{J}\mathbf{W}^*\mathbf{x}_t^*$ for some diagonal matrix $\mathbf{J}$ which has ± 1 on its diagonal. The affine equivariance thus means that the recovered components do not depend on the mixing matrix, except for their signs.

In the proposed SSA approach, the eigenvalues of $\mathbf{M}$ matrices provide a way of ordering nonstationary components (or pseudo-eigenvalues in the case of SSAcomb). Thus if the eigenvalues corresponding to the nonstationarity components are unique then $\mathbf{U}_n$ is well defined and we have $\mathbf{W}_n\mathbf{x}_t = \mathbf{J}\mathbf{W}_n^*\mathbf{x}_t^*$. However, for the stationary part all (pseudo-)eigenvalues are equal and thus $\mathbf{U}_s$ is not well defined and we actually have $\mathbf{W}_s\mathbf{x}_t = \mathbf{O}\mathbf{W}_n^*\mathbf{x}_t^*$ for some orthogonal matrix $\mathbf{O}$. However, as for finite data the eigenvalues are usually all distinct, the finite sample version will be affine equivariant. Note however that $\mathbf{W}_s\mathbf{x}_t$ is not necessarily equivalent to $\mathbf{s}_t$ and $\mathbf{W}_n\mathbf{x}_t$ is not necessarily equivalent to $\mathbf{n}_t$ as both $\mathbf{s}_t$ and $\mathbf{n}_t$ have the ambiguities mentioned above and only the two subspaces are well defined. This is quite similar to a non-Gaussian subspace analysis (NGCA) [39] - a BSS method for iid data that divides data into a Gaussian part and a non-Gaussian part. See also the supplementary material of [8] for a detailed discussion.

The aim of SSA is to separate the nonstationary and stationary subspaces, however, it is often desirable to have models where the interesting components are well defined (up to some minor identifiability issues such as their signs). In this work we do not consider a full analysis of the identifiability issues, just mention some possibilities. The main

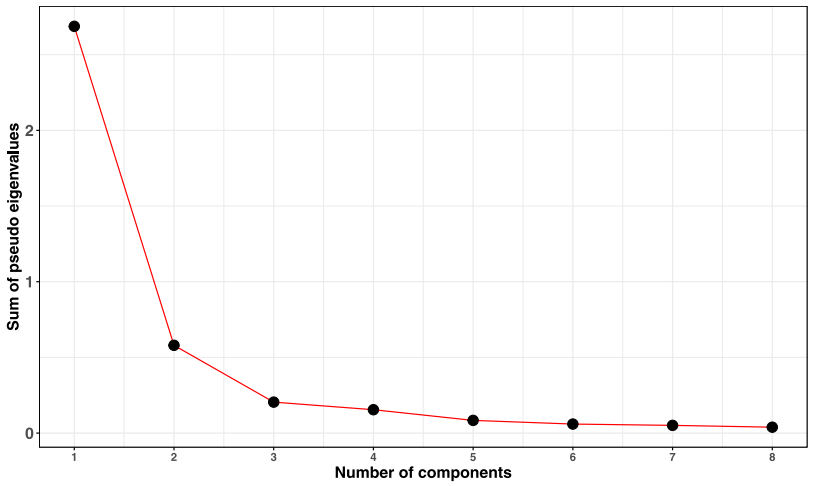

Fig. 3. Screeplot based on the column sums of the pseudo-eigenvalues from Table 1.

Table 1
Pseudo-eigenvalues of SSAcomb for a simulated time series.

	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8
$\mathbf{M}_m$	0.00	0.53	0.00	0.02	0.01	0.00	0.00	0.00
$\mathbf{M}_v$	2.67	0.03	0.03	0.07	0.03	0.03	0.02	0.02
$\mathbf{M}_\tau$	0.02	0.02	0.18	0.07	0.04	0.03	0.02	0.02

strategy we suggest in this context is to perform first SSA and then, if stronger assumptions on the stationary and/or non-stationary components can be made, use appropriate BSS methods and apply them to the recovered subspace. For general overviews of BSS models and assumptions they make on interesting components, see for example [3–7]. Using additional assumptions, stationary components are identifiable when the components follow for example a second-order source separation (SOS) model or a stationary independent time series model. For details, see for example [7]. In case of stationary time series methods such as AMUSE [40,41], SOBI [1,32,42], FixNA [43], gSOBI [33] or gJADE [44] can be applied to the identified stationary subspace. For the nonstationary subspace at least the following assumptions allow the estimation of the components. If all nonstationary components are independent and the marginal distributions are non-Gaussian, independent component analysis (ICA) methods can be used. For an overview of several ICA methods, see [4–6]. Another possibility is to assume a nonstationary source separation (NSS) model for the nonstationary components, that is, to assume that the mean is stationary, but each component has nonstationary variances which change in different patterns. This is for example a reasonable assumption for audio data. Interestingly, most NSS methods also divide the time series into intervals and jointly diagonalize statistics based upon them. See for example [2,45,46] for more details on NSS methods. In this context SSAsave seems to be a natural SSA method and could be for example compared to an approach where one first applies an NSS method and then tests each component for stationarity.

So far we have assumed that the dimensions of the two subspaces are known. In [8,11] a sequential likelihood ratio test for testing the dimension were proposed, and in [20] a similar approach was considered for a frequency domain approach. As similar inferential tools are so far missing for SSAcomb, we study next if the screeplot of the (pseudo-)eigenvalues give some indication of the choice of k. We apply SSAcomb based on six equal sized intervals for an eight-variate time series with $T = 4000$. The time series are simulated according to Setting 4 in the following section.

Table 1 contains the pseudo-eigenvalues and the screeplot shown in Fig. 3 is based on their column sums. The screeplot suggests three or four nonstationary components. Taking then into account the underlying pseudo-eigenvalues it is clear that the first component is nonstationary according to $\mathbf{M}_v$ only indicating nonstationarity in variance. The second component is nonstationary according to $\mathbf{M}_m$ and the third component is nonstationary according to $\mathbf{M}_\tau$. The fourth component has small values for all three $\mathbf{M}$-matrices. Therefore we conclude that the component may actually not be nonstationary. Notice that the true value of k is indeed three in this case.

4. Simulations and an example

In this section we compare the methods discussed above using simulation studies. The method that combines the search for all three types of nonstationarities is also illustrated by an example.

4.1. Simulations

As pointed out above, it depends on the purpose of the analysis whether the stationary subspace or the nonstationary subspace is of interest. The goal in this simulation study is therefore to evaluate how well the different methods can estimate the two subspaces under the assumption that the dimension of the subspace is known, as it is assumed in most of the SSA references mentioned above. Thus, based on the sample eigenvalues of the methods, the unmixing matrix estimates can be decomposed into $\hat{\mathbf{W}}_s$ and $\hat{\mathbf{W}}_n$ with dimensions $(p-k)\times p$ and $k\times p$, respectively.

For comparing the methods, we need a performance measure that takes into account the fact that the subspaces are only specified up to rotations. We therefore compute the projection matrices

$$\hat{\mathbf{P}}_s = \hat{\mathbf{W}}_s(\hat{\mathbf{W}}_s{}^\top\hat{\mathbf{W}}_s)^{-1}\hat{\mathbf{W}}_s{}^\top \quad \text{and} \quad \hat{\mathbf{P}}_n = \hat{\mathbf{W}}_n(\hat{\mathbf{W}}_n{}^\top\hat{\mathbf{W}}_n)^{-1}\hat{\mathbf{W}}_n{}^\top$$

and will compare them with the corresponding true projection matrices $\mathbf{P}_s$ and $\mathbf{P}_n$ by computing the squared distances

$$D_s^2 = \frac{1}{2}\|\mathbf{P}_s - \hat{\mathbf{P}}_s\|^2 \quad \text{and} \quad D_n^2 = \frac{1}{2}\|\mathbf{P}_n - \hat{\mathbf{P}}_n\|^2,$$

where $\|\cdot\|$ is again the Frobenius norm. As shown in [47,48], the performance indices can take values in $[0, \min\{k, p-k\}]$, with zero indicating a perfect recovery of the subspace.

All following simulations were done with R [49] using the packages JADE [36], LDRtools [50] and ssaBSS [51]. In all simulation settings the observed time points are $1,\dots,T$, where T varies from 1000 to 32000, and $p = 8$ with five stationary components $s_{1,t},\dots,s_{5,t}$ and three nonstationary components $n_{1,t}$, $n_{2,t}$ and $n_{3,t}$. Most components are based on moving average (MA) processes, autoregressive (AR) processes and autoregressive moving average (ARMA) processes. However, inspired by [15] we consider also time varying variance (TV-VAR) processes x_t with parameters α and β of the form $x_t = \tilde{h}_t\epsilon_t$ where $\tilde{h}_t^2 = h_t^2 + \alpha x_{t-1}^2$ where ϵ are iid $N(0,1)$, $x_0 = 0$ and $h_t = 10 - 10\sin(\beta\pi\ t/T + \pi/6)(1 + t/T)$. Similarly we consider also time varying autoregressive processes of order 1 (TV-AR1) where $x_t = a_t x_{t-1} + \epsilon_t$ with $x_0 = 0$, ϵ_t is iid $N(0,\sigma^2)$ and $a_t = 0.5\cos(2\pi\ t/T)$.

The four 8-variate settings under consideration are then:

Setting 1: $s_{1,t}$ is a MA(0.72, 0.24) process, $s_{2,t}$ is an AR(0.34, 0.27, 0.18) process, $s_{3,t}$ is an ARMA(0.34, 0.27, 0.18; 0.72, 0.15) process, $s_{4,t}$ is an AR(0.11, 0.58) process and $s_{5,t}$ is an MA(0.78) process. $n_{1,t} = y_t + \mu_t$ where y_t is an AR(0.7) process and $\mu_t = -1.52$ if $t \le \lfloor T/2\rfloor$ and otherwise 1.38. $n_{2,t} = y_t + \mu_t$ where y_t is an AR(0.5) process and $\mu_t = -0.75$ if $t \le \lfloor T/3\rfloor$, 0.84 if $t \in [\lfloor T/3\rfloor + 1, 2\lfloor T/3\rfloor]$ and otherwise -0.45.

Setting 2: $s_{1,t}$ is a MA(0.72) process, $s_{2,t}$ is a MA(0.34) process, $s_{3,t}$ is a MA(0.72, 0.15), $s_{4,t}$ is an MA(0.11, 0.58) process and $s_{5,t}$ is a MA(0.34, 0.27, 0.18) process. $n_{1,t}$ is TV-VAR(0.2,0.5), $n_{2,t}$ is TV-VAR(0.1,1) and $n_{3,t}$ is TV-VAR(0.05,0.01).

Setting 3: $s_{1,t}$ is an ARMA(0.14, 0.45; 0.72, 0.24) process, $s_{2,t}$ is an AR(0.34, 0.27, 0.18) process, $s_{3,t}$ is an ARMA(0.34, 0.27, 0.18; 0.72, 0.15) process, $s_{4,t}$ is an AR(0.11,0.58) process and $s_{5,t}$ is an AR(0.1, 0.1, 0.1, 0.1, 0.1) process. $n_{1,t}$ is a TV-AR process with $\sigma^2 = 0.8649$, $n_{2,t}$ consists of three independent blocks, i.e., for $t = 1$ to $t = \lfloor T/3\rfloor$ an AR(0.5) process with $\sigma^2 = 1$, for $t = \lfloor T/3\rfloor + 1$ to $t = 2\lfloor T/3\rfloor$ an AR(0.2) process with $\sigma^2 = 1.6384$ and for $t = 2\lfloor T/3\rfloor + 1$ to $t = T$ an AR(0.8) process with $\sigma^2 = 0.2304$. $n_{3,t}$ consists of two independent blocks, i.e., for times $t = 1$ to $t = \lfloor T/2\rfloor$ a MA(0.5) process with $\sigma^2 = 1$ and for $t = 2\lfloor T/2\rfloor + 1$ to $t = T$ a MA(0.9, 0.17) process with $\sigma^2 = 0.4624$

Setting 4: Has the same stationary components as used in Setting 3 and $n_{1,t}$ corresponds to $n_{1,t}$ in Setting 1, $n_{2,t}$ corresponds to $n_{2,t}$ in Setting 2 and $n_{3,t}$ corresponds to $n_{1,t}$ in Setting 3.

Thus, in Setting 1 the three nonstationary components have all nonstationary means, while in Setting 2 the means are all stationary but the variances are nonstationary. In Setting 3 the nonstationary information is in the correlation structure. Setting 4 uses then one nonstationary component from each of the previous settings meaning that here combining different approaches seems especially important. Examples of nonstationary components for Settings 1–3 with $T = 2000$ are shown in Figs. 4–6 showing that the violations of nonstationarity can take quite different forms.

In the simulation study we simulated 2000 source sets with various sample sizes T from all four settings and always used a random orthogonal mixing matrix as $\mathbf{A}$. Then for all methods we computed the distances between the true projections and the estimated projections when using $K = 2, 6$ and 12 equal-sized intervals. The use of $K > 12$ did not seem to improve the performance. In SSAcor and SSAcomb we used $\tau = 1$. Notice that in the spirit of SIR, the number of slices for SSAsir should exceed $k = 3$ and therefore SSAsir is not expected to work when $K = 2$.

The average distances are shown in Figs. 7–10. It is clearly visible that having only two intervals, that is $K = 2$, is always the poorest choice. Although this was an expected result for SSAsir, it was not so obvious for the other methods. The differences between $K = 6$ and $K = 12$ decrease with increasing sample size. However for smaller sample sizes the choice $K = 6$ is preferable. This confirms results familiar from NSS and SDR studies which show that it is important that the intervals contain enough observations to estimate the quantities of interest well enough.

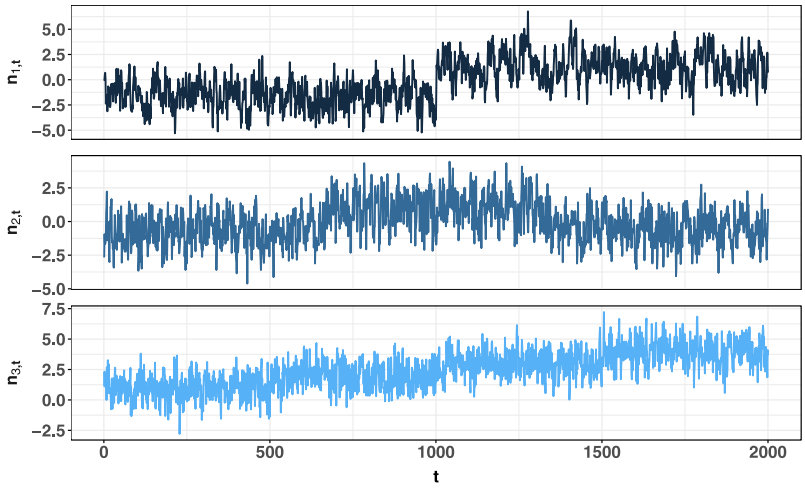

Fig. 4. Example of the nonstationary components in Setting 1 for $T = 2000$.

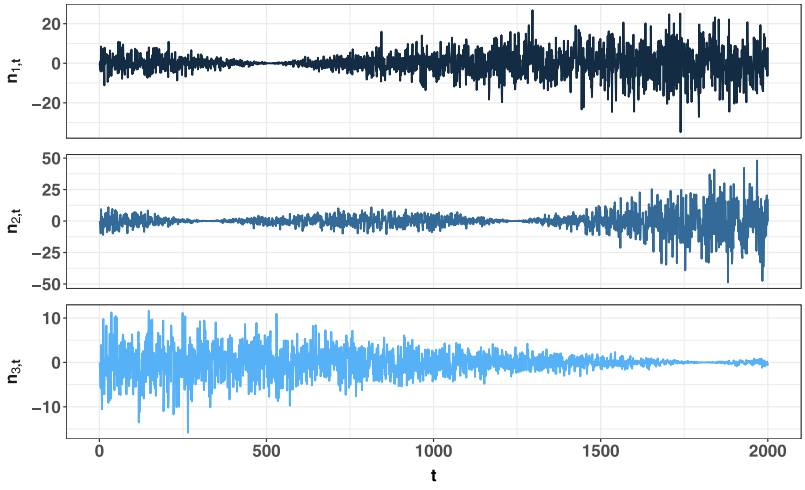

Fig. 5. Example of the nonstationary components in Setting 2 for $T = 2000$.

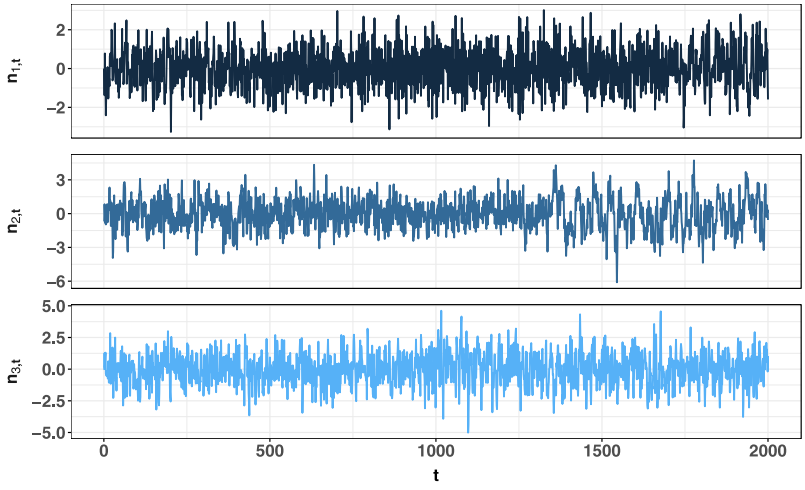

Fig. 6. Example of the nonstationary components in Setting 3 for $T = 2000$.

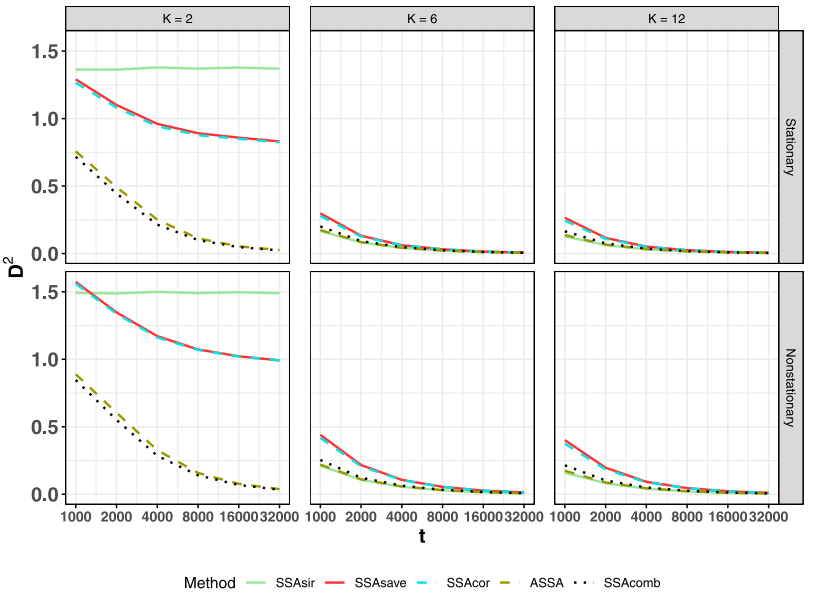

Fig. 7. Average squared distances between the true and estimated subspaces in Setting 1.

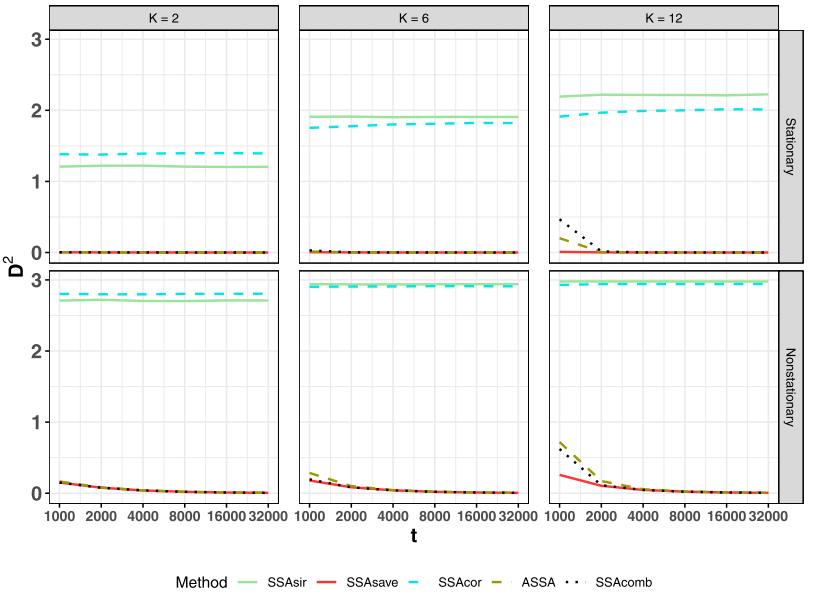

Fig. 8. Average squared distances between the true and estimated subspaces in Setting 2.

As seen in Fig. 7, in case of Setting 1 all methods are able to estimate the subspaces when there are enough intervals. This is a bit surprising result for SSAcor. From the results based on Setting 2 in Fig. 8 it is clear that SSAsir and SSAcor do not work at all. ASSA and SSAcomb work well despite containing matrices which, when used by themselves, fail completely. The results based on Setting 3 in Fig. 9 indicate that SSAcor and SSAcomb perform clearly best while ASSA does worse than SSAsave and is therefore affected by the poor performance of SSAsir. Finally, as seen in Fig. 10, in Setting 4, where all different types of nonstationarities are present, it is clear that SSAcomb is the only method that can estimate the subspaces properly. It requires however a much larger sample size to detect all three nonstationary components. ASSA seems to be constantly missing one nonstationary component which is not surprising as it is not looking for nonstationarity in

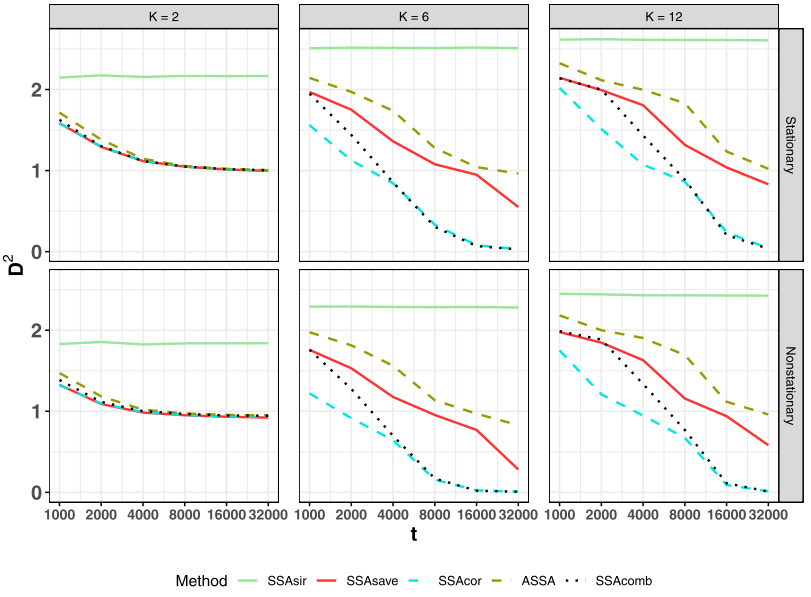

Fig. 9. Average squared distances between the true and estimated subspaces in Setting 3.

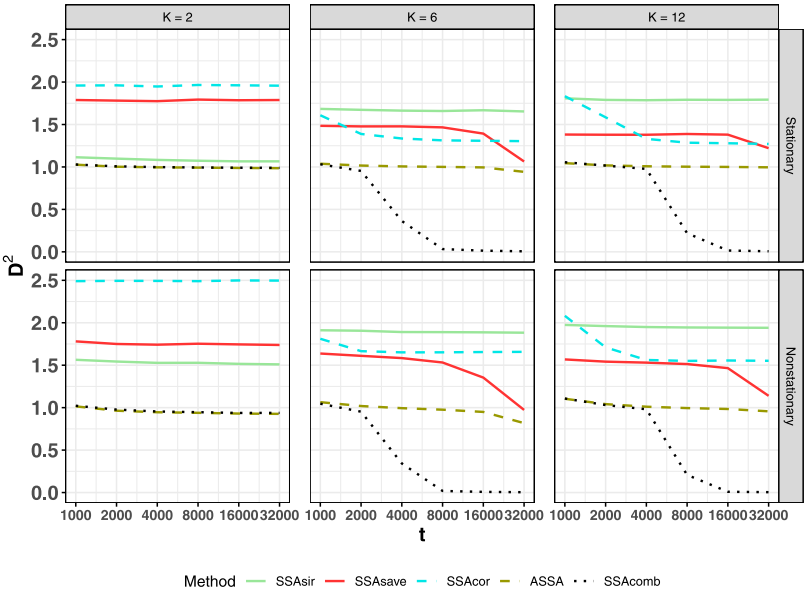

Fig. 10. Average squared distances between the true and estimated subspaces in Setting 4.

autocorrelation structures. To study the effect of dimensionality on performances we increased the number of components in Settings 1–4 from $p = 8$ to $p = 20$ so that we added 12 additional stationary components (see Appendix C for details). The results shown in Figs. C.17–C.20 in Appendix C indicate that especially in Settings 1, 3 and 4 the proposed SSA methods perform well only if the time series length is high enough. Again, the choice $K = 6$ seems to be preferable over the two other options. Finally, to compare the performances of methods to that of DSSA introduced in [20] we performed a small simulation study where DSSA was used to recover the stationary subspace in Settings 1–4. As the computation times for DSSA were much higher than for the other methods (see Fig. C.22 in Appendix C), we only simulated 50 source sets with various sample sizes T from all four settings. The results shown in Fig. C.21 in Appendix C illustrate that DSSA outperforms

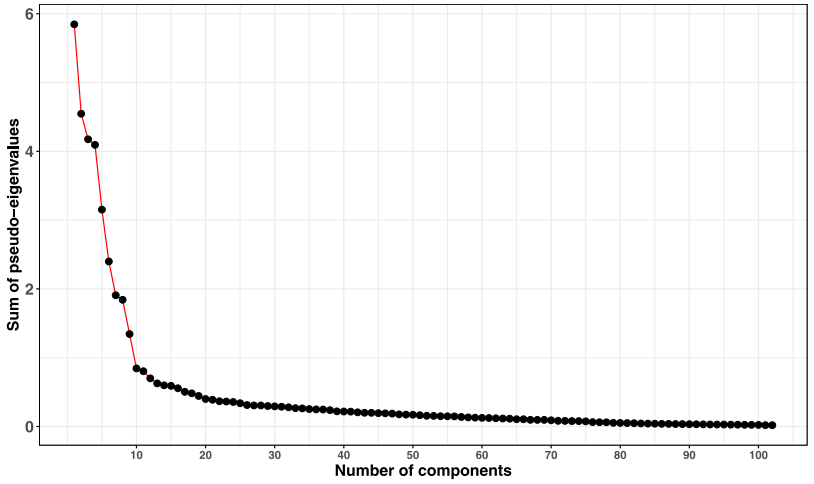

Fig. 11. Screeplot of the sums of the pseudo-eigenvalues based on MEG data.

Table 2
Pseudo-eigenvalues of SSAcomb based on the MEG data.

	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	d_9	d_{10}
$\mathbf{M}_m$	0.08	0.06	0.00	0.01	0.01	0.41	0.00	0.00	0.03	0.00
$\mathbf{M}_v$	2.90	2.26	2.47	2.07	1.59	1.03	1.89	1.18	0.68	0.45
$\mathbf{M}_\tau$	2.87	2.22	1.70	2.02	1.55	0.96	0.02	0.66	0.64	0.39

its natural competitor SSAcomb only in Setting 2 and with very small sample sizes. Interestingly, in Setting 4, where all three types of nonstationarities are encountered, DSSA performs equally well with SSAsave and SSAcor and worse than ASSA and SSAcomb.

Thus, to conclude, based on the simulation studies SSAcomb seems to be a preferable method as it works quite well independently of the nature of the underlying stationarity violations. However, large sample sizes are preferable and the number of intervals K plays a role in the performance. In the simulation studies above, a moderate number of $K = 6$ seems as a reasonable choice.

4.2. Example

We applied the proposed SSAcomb method to a magnetoencephalographic (MEG) data of length $T = 221710$ that was recorder using $p = 102$ magnetometers at the Center for Interdisciplinary Brain Research, Department of Psychology, University of Jyväskylä. Magnetometers record brain activity indirectly by measuring changes in magnetic fields generated by electrical currents occurring in the brain. As the measurements are taken from top of the head, it is realistic to assume that the observed signals are mixtures of actual brain activity signals. The signals are also mixed with artifacts that occur due to external physiological factors such as moving the head, tensing muscles and blinking and moving eyes. Such mixing is also seen in Fig. B.15 in Appendix B where we plot ten first MEG-signals. The goal of the analysis is then to recover the brain activity of interest from the observed mixture. Notice that KL-SSA was applied to EEG data analysis with the similar goal in mind in [8,9].

Our interest was to see if SSAcomb can recover nonstationary components that possibly correspond to interesting brain activity or physiological signals. As the length of the data was pretty large, we applied SSAcomb with $K = 12$, and plotted the screeplot of the sums of the pseudo-eigenvalues in Fig. 11.

Based on the screeplot we note that the method is able to detect ten nonstationary components. We then list the ten pseudo-eigenvalues in Table 2 and conclude that the components 1–5 and 8–10 were detected due to nonstationarity in variance and in autocorrelation. Component 6 has all three types of nonstationarities and component 7 is nonstationary with respect to the variance.

Fig. 12 plots the ten nonstationary components recovered by SSAcomb.

Due to large number of observations, nonstationarities in autocorrelation are difficult to observe, but nonstationarities in mean and variance are clearly visible in Fig. 12. In Fig. B.16 in Appendix B we show the topography plots related to each signal. The topography plots demonstrate which part of the brain is activated with dark colors (red and blue) indicating high activity. Such plots help to detect those components that are related to the artifacts and those components that are

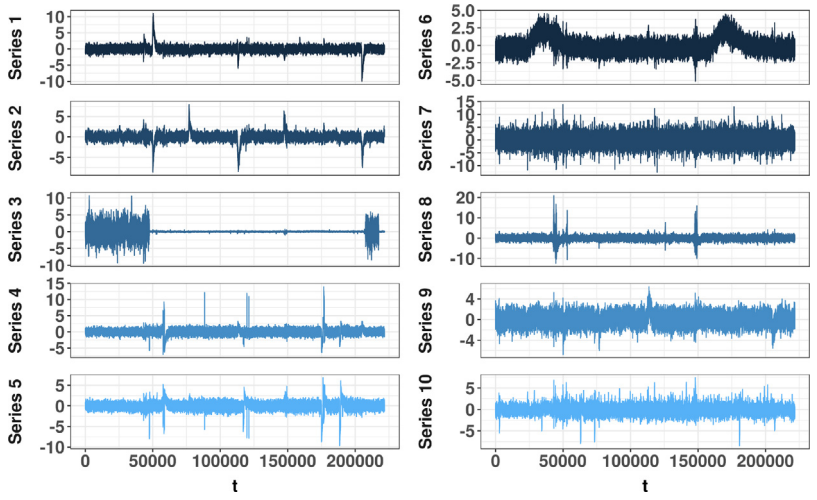

Fig. 12. Ten nonstationary components recovered by SSAcomb method with $K = 12$.

related to the brain activity that is relevant to the study at hand. We emphasize that the in-depth interpretation of the topography plots and detected components should be done by subject experts. Notice that many MEG studies proceed by removing artifacts before trying to detect the brain activity signals of interest, see for example [52] and references therein.

5. Conclusion

In this paper we gave a statistical perspective on stationary subspace analysis. We showed how SSA can be formulated as a supervised dimension reduction method, and in that spirit, proposed methods that can detect nonstationarities in mean, variance, autocorrelation, and in all of them. The method thus extends the classical analytic SSA into a more general time series setting. Simulation studies were used to study the effect of number of intervals K and data dimensionality p on performance under various simulation settings. Based on simulation results in Section 4 it was noted that the SSAcomb method which combined all three nonstationarity measures and used a moderate number of intervals was among the best performing methods in all settings. The method can thus be seen as a safe choice unless some additional information on latent sources are available. If latent sources are expected to exhibit some special type of nonstationarity, then it is beneficial to use the method that is developed for that specific purpose. Another option would be to weight the matrices in SSAcomb to put more emphasis on some special structures. As shown in Section 2.6, if the time series are known to have some seasonal of cyclic components, then we advice selecting unequal sized interval widths to make sure that such nonstationarities can be detected.

In this study we assumed that the number of nonstationary components, k, is known. However, most often this parameter value is unknown and needs to be estimated. In [8,11] a sequential likelihood ratio test was proposed for determining the dimension of the stationary subspace in iid case. In the context of ASSA, [19] mentioned that eigenvalues can give guidance for choosing the dimension, but they did not pursue this idea any further. When using SSAsir, SSAsave, SSAcor or SSAcomb, a possible test for testing the null hypothesis $H_0 : \ k = k_0$, where k is the true dimension of the nonstationary subspace and k_0 is the proposed dimension, could also be based on the eigenvalues of matrices $\mathbf{M}$ and resampling-based procedures in a similar fashion as in [31,53], for example. An estimate for the nonstationary subspace dimension could then be based on the sequential testing as in [54]. We leave this for a future work. In this context we will also investigate if the detection of the type of nonstationarity can be formalized and how the SSA methods can help to detect change points in the multivariate time series. To tackle these issues we will study in future the asymptotic properties of the methods discussed in this paper in more detail.

Data availability

The authors do not have permission to share data.

Acknowledgments

The work of KN was supported by the Austrian Science Fund (FWF) under grant P31881-N32. We thank Professor Jarmo Hämäläinen for providing us the MEG data and the reviewers for their insightful comments which helped to improve the paper.

Appendix A. Settings and a ddditional figures comparing SSAsir, SSAsave and SSAcor

The settings used in Section 2.6 are

Fig. 1 (left): $x_t = y_t + \mu_t$ where y_t is an AR(0.7) process and $\mu_t = -1.52$ if $t \leq \lfloor T/2 \rfloor$ and otherwise 1.38.

Fig. 1 (right): $x_t = y_t + \sin(t/150\pi)$ where y_t is an AR(0.2) process with $\sigma^2 = 0.25$.

Fig. 2 (left): x_t is TV-VAR(0.8, 1) process.

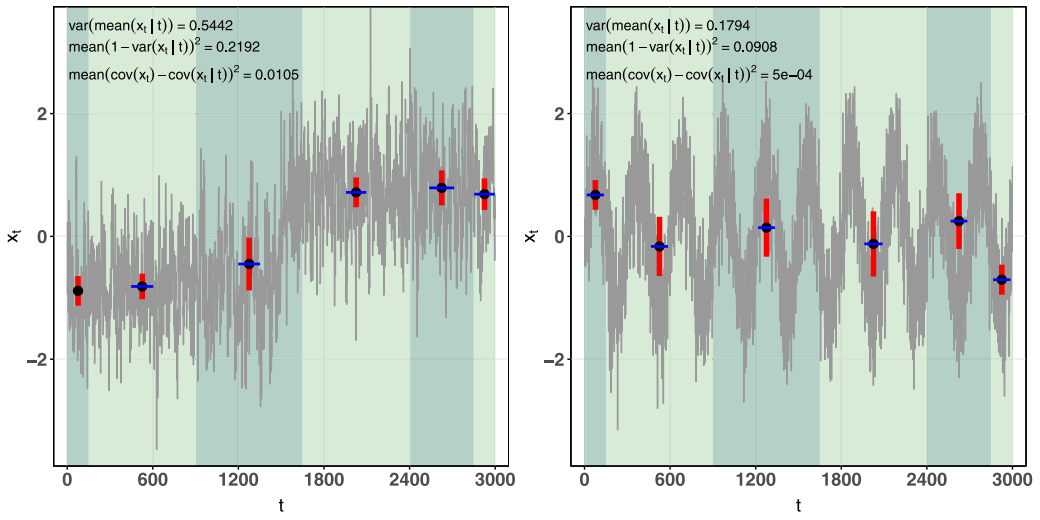

Fig. A.13. Visualization of SSAsir, SSAsave and SSAcor for a component with nonstationarity in mean (left panel trend, right panel seasonality) when the six intervals have different sizes. The black dots represent the interval mean, the height of vertical red bar the interval variance and the width of blue horizontal bar the interval autocovariance.

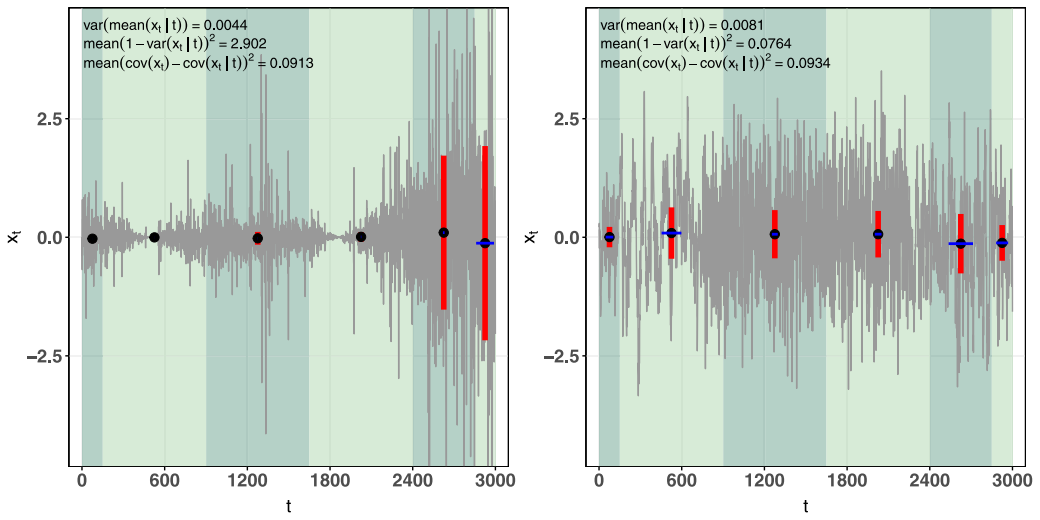

Fig. A.14. Visualization of SSAsir, SSAsave and SSAcor for a component with nonstationarity in variance (left panel) and in autocorrelation (right panel) when the six intervals have different sizes. The black dots represent the interval mean, the height of vertical red bar the interval variance and the width of blue horizontal bar the interval autocovariance.

Fig. 2 (right): x_t consists of three independent blocks, i.e., for $t = 1$ to $t = \lfloor T/3 \rfloor$ an AR(0.9) process with $\sigma^2 = 0.25$, for $t = \lfloor T/3 \rfloor + 1$ to $t = 2\lfloor T/3 \rfloor$ an AR(0.2) process with $\sigma^2 = 0.09$ and for $t = 2\lfloor T/3 \rfloor + 1$ to $t = T$ an AR(0.8) process with $\sigma^2 = 2.25$.

Visualizations of SSAsir, SSAsave and SSAcor for the same series as in Section 2.6 when the intervals are of unequal lengths.

Appendix B. Additional figures for the MEG example

See Figs. B.15 and B.16.

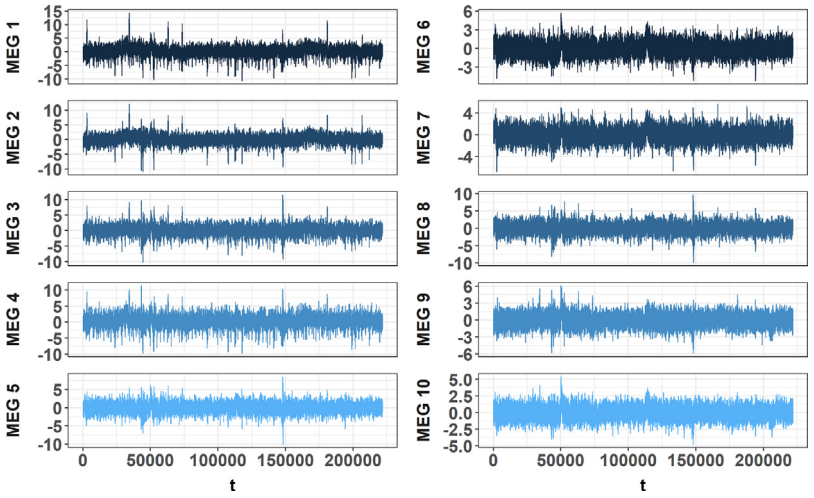

Fig. B.15. Ten observed MEG signals.

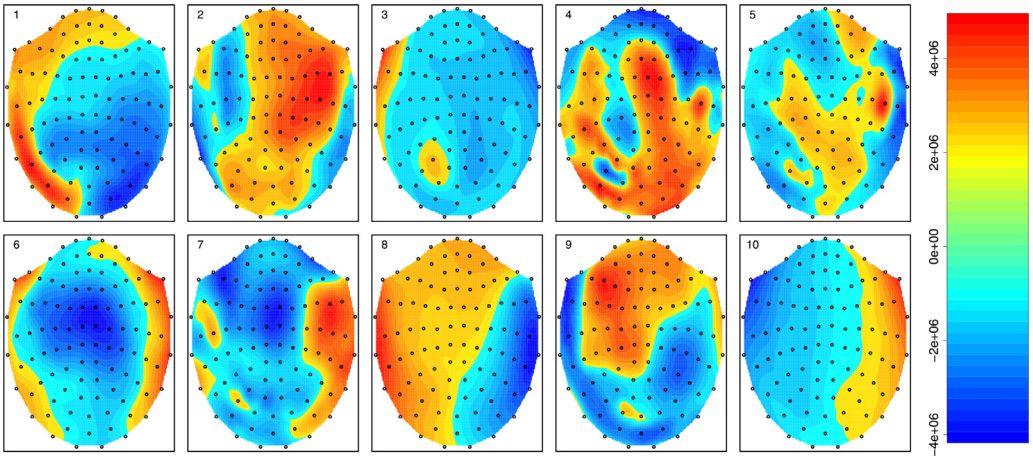

Fig. B.16. Topography plots based on ten nonstationary components recovered from the MEG data. The topographies illustrate which part of the brain is activated related to each of the nonstationary components. Here dark red and dark blue indicate high activation.

Appendix C. Additional simulation results

We repeated the simulations mentioned in Section 4 using $p = 20$ time series instead of $p = 8$. Each setting included the same non-stationary components as before. All settings used the same 12 stationary components. All stationary components in the original settings were used and, in addition, MA(0.5), MA(−0.4), MA(−0.28) and MA(0.34) processes were added. Results are presented in Figs. C.17–C.20.

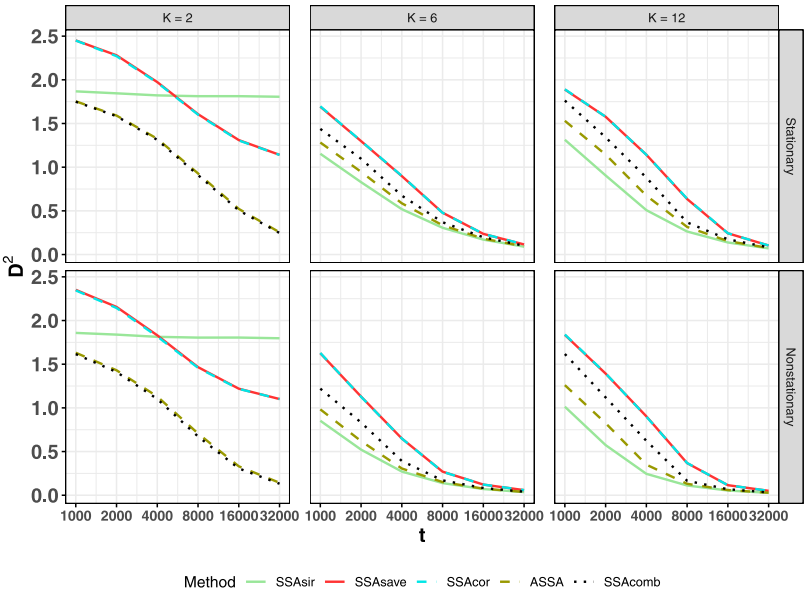

Fig. C.17. Average squared distances between the true and estimated subspaces in Setting 1 with $p = 20$.

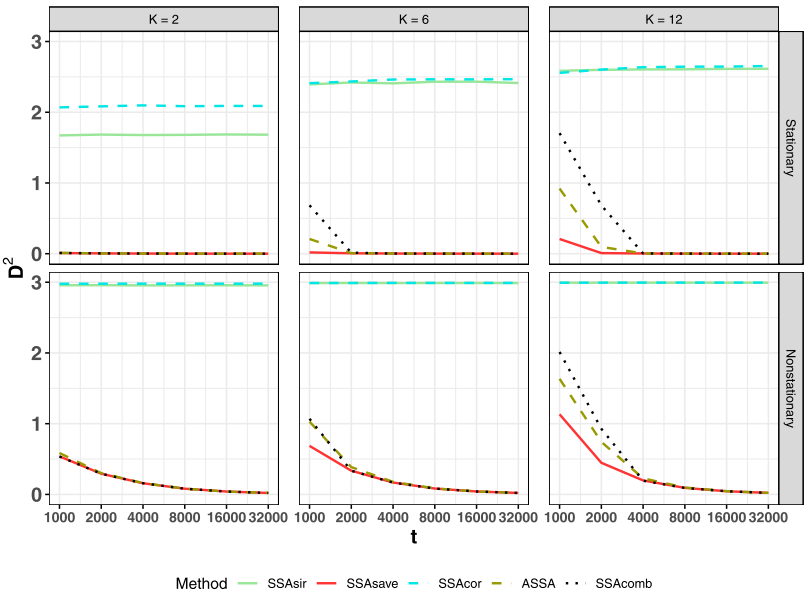

Fig. C.18. Average squared distances between the true and estimated subspaces in Setting 2 with $p = 20$.

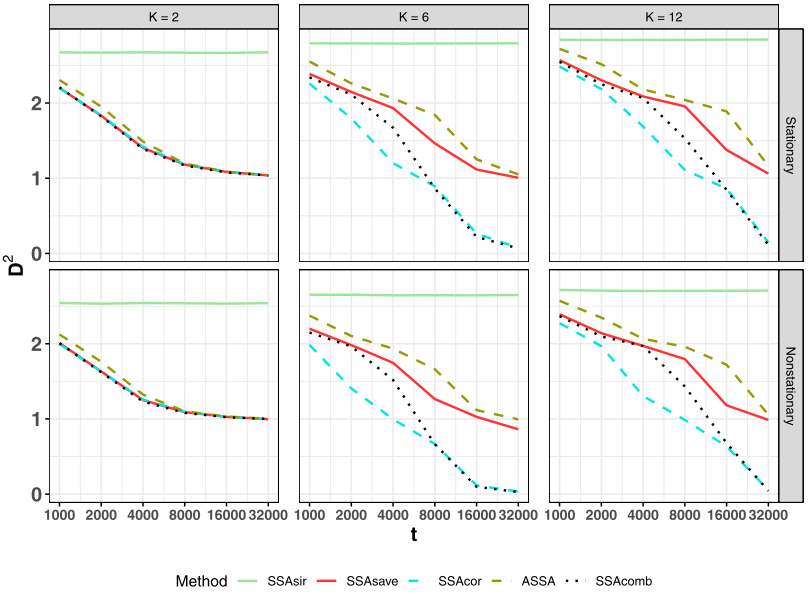

Fig. C.19. Average squared distances between the true and estimated subspaces in Setting 3 with $p = 20$.

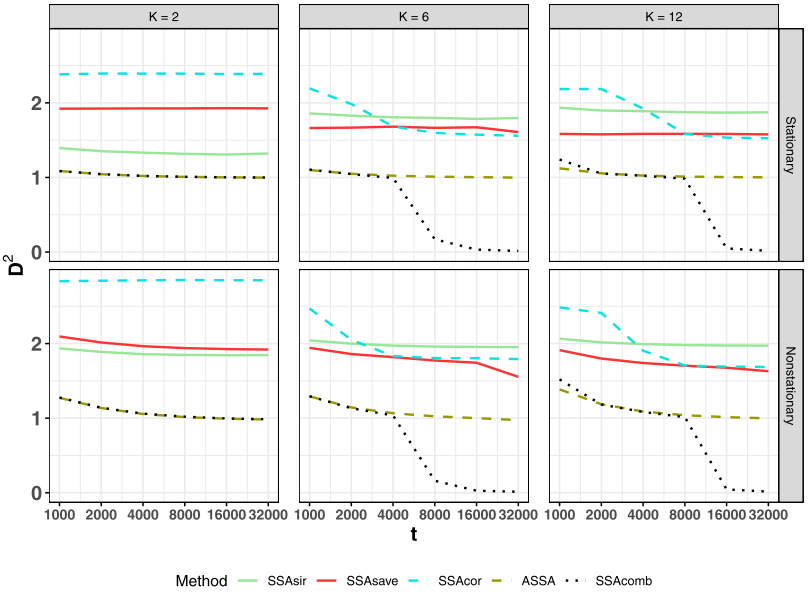

Fig. C.20. Average squared distances between the true and estimated subspaces in Setting 4 with $p = 20$.

We also added DSSA as one of the methods and repeated the simulations mentioned in Section 4. As DSSA was used with time series of length under 1000 [20], we also added $T = 500$ as one of the time series lengths. Fig. C.21 presents the results for $K = 6$, as it was deemed as the safest choice. As DSSA was significantly slower as compared to the other methods (Fig. C.22), we used only 50 simulated source sets. As DSSA performed the best in Setting 2, the simulations for this setting were repeated with 1000 source sets in order to validate the results. The code used for DSSA can be found in https://github.com/tambstat/SSA/blob/main/dssa.R.

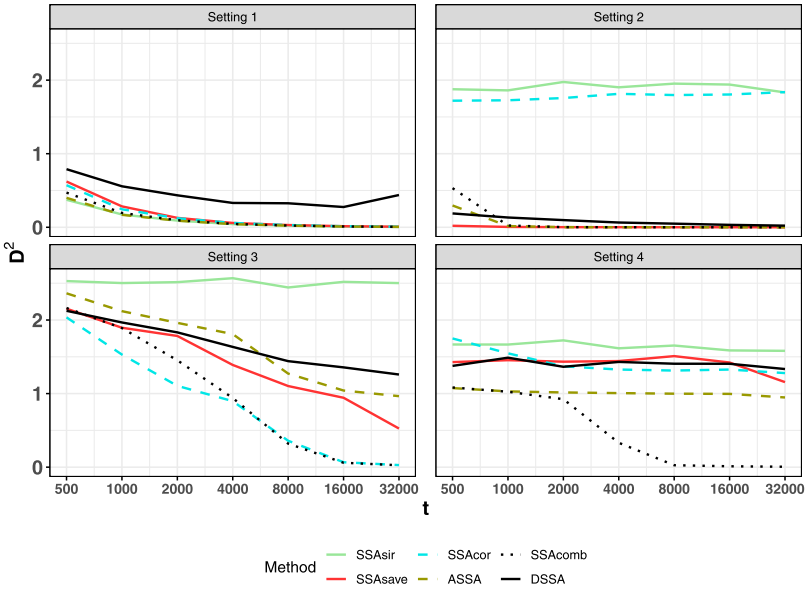

Fig. C.21. Average squared distances between the true and estimated stationary subspace in all simulation settings.

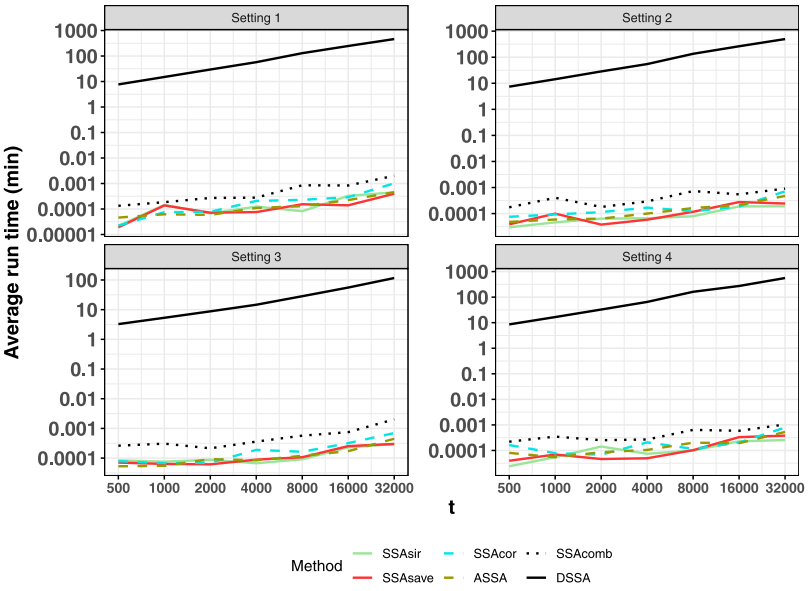

Fig. C.22. Average run times for all methods (in minutes) and simulation settings.

References

[1] J. Miettinen, K. Illner, K. Nordhausen, H. Oja, S. Taskinen, F. Theis, Separation of uncorrelated stationary time series using autocovariance matrices, J. Time Series Anal. 37 (2016) 337–354.

[2] S. Choi, A. Cichocki, Blind separation of nonstationary sources in noisy mixtures, Electron. Lett. 36 (2000) 848–849.

[3] A. Cichocki, S.-I. Amari, Adaptive Blind Signal and Image Processing: Learning Algorithms and Applications, John Wiley & Sons, New York, 2002, http://dx.doi.org/10.1002/0470845899.

[4] P. Comon, C. Jutten, Handbook of Blind Source Separation: Independent Component Analysis and Applications, Academic Press, Amsterdam, 2010.
[5] T. Adali, M. Anderson, G.-S. Fu, Diversity in independent component and vector analyses: identifiability, algorithms, and applications in medical imaging, IEEE Signal Process. Mag. 31 (2014) 18–33, http://dx.doi.org/10.1109/MSP.2014.2300511.
[6] K. Nordhausen, H. Oja, Independent component analysis: A statistical perspective, WIREs: Comput. Stat. 10 (2018) e1440, http://dx.doi.org/10.1002/wics.1440.
[7] Y. Pan, M. Matilainen, S. Taskinen, K. Nordhausen, A review of second-order blind identification methods, WIREs Comput. Stat. 14 (2022) e1550, http://dx.doi.org/10.1002/wics.1550.
[8] P. von Bünau, F.C. Meinecke, F.C. Király, K.-R. Müller, Finding stationary subspaces in multivariate time series, Phys. Rev. Lett. 103 (2009) 214101, http://dx.doi.org/10.1103/PhysRevLett.103.214101.
[9] P. von Bünau, F.C. Meinecke, S. Scholler, K. Müller, Finding stationary brain sources in EEG data, in: 2010 Annual International Conference of the IEEE Engineering in Medicine and Biology, 2010, pp. 2810–2813, http://dx.doi.org/10.1109/IEMBS.2010.5626537.
[10] F.C. Meinecke, P. von Bünau, M. Kawanabe, K. Müller, Learning invariances with stationary subspace analysis, in: 2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops, 2009, pp. 87–92, http://dx.doi.org/10.1109/ICCVW.2009.5457715.
[11] D.A.J. Blythe, P. von Bunau, F.C. Meinecke, K. Müller, Feature extraction for change-point detection using stationary subspace analysis, IEEE Trans. Neural Netw. Learn. Syst. 23 (4) (2012) 631–643, http://dx.doi.org/10.1109/TNNLS.2012.2185811.
[12] D.A.J. Blythe, F.C. Meinecke, P. von Bünau, K.-R. Müller, Explorative data analysis for changes in neural activity, J. Neural Eng. 10 (2013) 026018, http://dx.doi.org/10.1088/1741-2560/10/2/026018.
[13] A. Aue, L. Horvath, Structural breaks in time series, J. Time Series Anal. 34 (1) (2013) 1–16, http://dx.doi.org/10.1111/j.1467-9892.2012.00819.x.
[14] R.S. Tsay, Analysis of Financial Time Series, second ed., in: Wiley Series in Probability and Statistics, Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, 2005, p. xxii+605, http://dx.doi.org/10.1002/0471746193.
[15] V. Patilea, H. Raissi, Testing second-order dynamics for autoregressive processes in presence of time-varying variance, J. Amer. Statist. Assoc. 109 (2014) 1099–1111, http://dx.doi.org/10.1080/01621459.2014.884504.
[16] A. Cardinali, G. Nason, Costationarity of locally stationary time series, J. Time Ser. Econom. 2 (2) (2010) 1–33.
[17] I. Horev, F. Yger, M. Sugiyama, Geometry-aware stationary subspace analysis, in: R.J. Durrant, K. Kim (Eds.), Proceedings of the 8th Asian Conference on Machine Learning, ACML 2016, Hamilton, New Zealand, November 16–18, 2016, in: JMLR Workshop and Conference Proceedings, vol. 63, JMLR.org, 2016, pp. 430–444.
[18] S. Kaltenstadler, S. Nakajima, K. Müller, W. Samek, Wasserstein stationary subspace analysis, IEEE J. Sel. Top. Sign. Proces. 12 (2018) 1213–1223, http://dx.doi.org/10.1109/JSTSP.2018.2873987.
[19] S. Hara, Y. Kawahara, T. Washio, P. von Bünau, Stationary subspace analysis as a generalized eigenvalue problem, in: K.W. Wong, B.S.U. Mendis, A. Bouzerdoum (Eds.), Neural Information Processing. Theory and Algorithms, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 422–429.
[20] R.R. Sundararajan, M. Pourahmadi, Stationary subspace analysis of nonstationary processes, J. Time Series Anal. 39 (2018) 338–355, http://dx.doi.org/10.1111/jtsa.12274.
[21] Y. Ma, L. Zhu, A review on dimension reduction, Inte. Stat. Rev. 81 (2013) 134–150.
[22] B. Li, Sufficient Dimension Reduction Methods and Applications with R, Chapman and Hall/CRC, Boca Raton, 2018.
[23] K.-C. Li, Sliced inverse regression for dimension reduction, J. Amer. Statist. Assoc. 86 (1991) 316–327.
[24] R. Cook, SAVE: a method for dimension reduction and graphics in regression, Comm. Statist. Theory Methods 29 (2000) 2109–2121.
[25] R.D. Cook, F. Critchley, Identifying regression outliers and mixtures graphically, J. Amer. Statist. Assoc. 95 (2000) 781–794.
[26] Z. Ye, R.E. Weiss, Using the bootstrap to select one of a new class of dimension reduction methods, J. Amer. Statist. Assoc. 98 (464) (2003) 968–979.
[27] L.-X. Zhu, M. Ohtaki, Y. Li, On hybrid methods of inverse regression-based algorithms, Comput. Statist. Data Anal. 51 (5) (2007) 2621–2635.
[28] A.J. Shaker, L.A. Prendergast, Iterative application of dimension reduction methods, Electron. J. Stat. 5 (2011) 1471–1494.
[29] M. Matilainen, C. Croux, K. Nordhausen, H. Oja, Supervised dimension reduction for multivariate time series, Econom. Stat. 4 (2017) 57–69, http://dx.doi.org/10.1016/j.ecosta.2017.04.002.
[30] M. Matilainen, C. Croux, K. Nordhausen, H. Oja, Sliced average variance estimation for multivariate time series, Statistics 53 (2019) 630–655, http://dx.doi.org/10.1080/02331888.2019.1605515.
[31] K. Nordhausen, H. Oja, D.E. Tyler, Asymptotic and bootstrap tests for subspace dimension, J. Multivariate Anal. 188 (2022) 104830, http://dx.doi.org/10.1016/j.jmva.2021.104830.
[32] A. Belouchrani, K. Abed Meraim, J.-F. Cardoso, E. Moulines, A blind source separation technique based on second order statistics, IEEE Trans. Signal Process. 45 (1997) 434–444.
[33] J. Miettinen, M. Matilainen, K. Nordhausen, S. Taskinen, Extracting conditionally heteroskedastic components using independent component analysis, J. Time Series Anal. 41 (2020) 293–311.
[34] P. Comon, Tensor diagonalization, a useful tool in signal processing, IFAC Proc. Vol. 27 (8) (1994) 77–82, IFAC Symposium on System Identification (SYSID'94), Copenhagen, Denmark, 4-6 July.
[35] D. Clarkson, A least squares version of algorithm AS 211: The F-G diagonalization algorithm, Appl. Stat. 37 (1988) 317–321.
[36] J. Miettinen, K. Nordhausen, S. Taskinen, Blind source separation based on joint diagonalization in R: the packages JADE and BSSasymp, J. Stat. Softw. 76 (2017) 1–31, http://dx.doi.org/10.18637/jss.v076.i02.
[37] N. Piccolotto, M. Bögl, T. Gschwandtner, C. Muehlmann, K. Nordhausen, P. Filzmoser, S. Miksch, TBSSvis: visual analytics for temporal blind source separation, Vis. Inform. 6 (4) (2022) 51–66, http://dx.doi.org/10.1016/j.visinf.2022.10.002.
[38] J. Miettinen, S. Taskinen, K. Nordhausen, H. Oja, Fourth moments and independent component analysis, Statist. Sci. 30 (2015) 372–390, http://dx.doi.org/10.1214/15-STS520.
[39] G. Blanchard, M. Kawanabe, M. Sugiyama, V. Spokoiny, K.-R. Müller, In search of non-gaussian components of a high-dimensional distribution, J. Mach. Learn. Res. 7 (2006) 247–282.
[40] L. Tong, V. Soon, Y. Huang, R. Liu, AMUSE: a new blind identification algorithm, in: Proceedings of IEEE International Symposium on Circuits and Systems, IEEE, 1990, pp. 1784–1787.
[41] J. Miettinen, K. Nordhausen, H. Oja, S. Taskinen, Statistical properties of a blind source separation estimator for stationary time series, Statist. Probab. Lett. 82 (2012) 1865–1873.
[42] S. Taskinen, J. Miettinen, K. Nordhausen, A more efficient second order blind identification method for separation of uncorrelated stationary time series, Statist. Probab. Lett. 116 (2016) 21–26, http://dx.doi.org/10.1016/j.spl.2016.04.007.
[43] Z. Shi, Z. Jiang, F. Zhou, A fixed-point algorithm for blind source separation with nonlinear autocorrelation, J. Comput. Appl. Math. 223 (2) (2009) 908–915, http://dx.doi.org/10.1016/j.cam.2008.03.009.
[44] M. Matilainen, K. Nordhausen, H. Oja, New independent component analysis tools for time series, Statist. Probab. Lett. 105 (2015) 80–87.
[45] D.-T. Pham, J.-F. Cardoso, Blind separation of instantaneous mixtures of nonstationary sources, IEEE Trans. Signal Process. 49 (9) (2001) 1837–1848, http://dx.doi.org/10.1109/78.942614.

[46] K. Nordhausen, On robustifying some second order blind source separation methods for nonstationary time series, Statist. Papers 55 (1) (2014) 141–156.
[47] L.J. Crone, D.S. Crosby, Statistical applications of a metric on subspaces to satellite meteorology, Technometrics 37 (1995) 324–328.
[48] E. Liski, K. Nordhausen, H. Oja, A. Ruiz-Gazen, Combining linear dimension reduction subspaces, in: C. Agostinelli, A. Basu, P. Filzmoser, D. Mukherjee (Eds.), Recent Advances in Robust Statistics: Theory and Applications, Springer India, New Delhi, 2016, pp. 131–149.
[49] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, 2022, URL: https://www.R-project.org/.
[50] E. Liski, K. Nordhausen, H. Oja, A. Ruiz-Gazen, LDRTools: tools for linear dimension reduction, 2018, URL: https://CRAN.R-project.org/package=LDRTools. R package version 0.2-1.
[51] M. Matilainen, L. Flumian, K. Nordhausen, S. Taskinen, ssaBSS: stationary subspace analysis, 2022, URL: https://CRAN.R-project.org/package=ssaBSS. R package version 0.1.1.
[52] J. Escudero, R. Hornero, D. Abasolo, A. Fernandez, M. Lopez-Coronado, Artifact removal in magnetoencephalogram background activity with independent component analysis, IEEE Trans. Biomed. Eng. 54 (11) (2007) 1965–1973, http://dx.doi.org/10.1109/TBME.2007.894968.
[53] K. Nordhausen, H. Oja, D.E. Tyler, J. Virta, Asymptotic and bootstrap tests for the dimension of the non-gaussian subspace, IEEE Signal Process. Lett. 24 (2017) 887–891, http://dx.doi.org/10.1109/LSP.2017.2696880.
[54] J. Virta, K. Nordhausen, Determining the signal dimension in second order source separation, Statist. Sinica 31 (2021) 135–156.