STATIONARY SUBSPACE ANALYSIS FOR SPATIAL DATA

Perttu Saarela
Department of Mathematics and Statistics
University of Helsinki
Finland
perttu.saarela@helsinki.fi

Klaus Nordhausen
Department of Mathematics and Statistics
University of Helsinki
Finland
klaus.nordhausen@helsinki.fi

Jaakko Pere
School of Business
Aalto University
Finland
jaakko.pere@aalto.fi

Anne M. Ruiz
Toulouse School of Economics
University Toulouse Capitole
France
anne.ruiz-gazen@tse-fr.eu

ABSTRACT

Stationary subspace analysis (SSA) is a blind source separation framework that decomposes linearly mixed multivariate data into stationary and nonstationary components. We extend SSA to spatially indexed data by introducing spatial stationary subspace analysis (spSSA), which explicitly accounts for spatial dependence. We propose three estimation procedures for the unmixing matrix based on first- and second-order spatial statistics. Each procedure targets a different type of nonstationarity and can be formulated as the solution to a generalized eigenvalue problem. To address situations where multiple forms of nonstationarity are present simultaneously, we combine the three procedures using approximate joint diagonalization. Simulation studies demonstrate that this combined approach yields superior separation performance. When the dimension of the nonstationary subspace is known, the proposed methods reliably recover the latent stationary and nonstationary components. However, determining this dimension remains a fundamental challenge in SSA, for which no generally accepted solution currently exists. Building on our estimation procedures, we propose a novel data augmentation approach to estimate the dimension of the nonstationary subspace and demonstrate its effectiveness through simulation studies. The proposed methodology is easily transferable to time series settings, making it of broader methodological interest.

***K*eywords** blind source separation · data augmentation · dimension reduction · multivariate spatial data · nonstationarity · order determination

1 Introduction

In spatial data analysis, observations $x(\mathbf{u}_i) \in \mathbb{R}$, $i = 1, \ldots, n$, are collected over a spatial domain $\mathcal{U} \subset \mathbb{R}^k$, where $\mathbf{u}_i \in \mathcal{U}$ denotes the sampling location. A central principle guiding spatial modeling is *Tobler's First Law of Geography* (Tobler, 1970), which states that "everything is related to everything else, but near things are more related than distant things." Consequently, observations recorded at nearby locations tend to be more similar than those farther apart, and this spatial dependence must be explicitly accounted for.

Spatial dependence is commonly characterized through the covariance function

$$C_x(\mathbf{u}_i, \mathbf{u}_j) = \mathrm{E}\big[\big(x(\mathbf{u}_i) - \mathrm{E}[x(\mathbf{u}_i)]\big)\big(x(\mathbf{u}_j) - \mathrm{E}[x(\mathbf{u}_j)]\big)\big] .$$

For more tractable modeling, spatial data are often assumed to arise from a weakly stationary square-integrable random field, implying a constant mean $\mathrm{E}[x(\mathbf{u})] = \mu$ and a covariance depending only on the separation vector $\mathbf{h} = \mathbf{u}_i - \mathbf{u}_j$. If the covariance depends solely on $h = \|\mathbf{h}\|$, the process is further said to be isotropic. Under these assumptions, several parametric covariance models exist, with the Matérn family of functions being among the most widely used (Guttorp and Gneiting, 2006).

In many applications, several variables are observed at each spatial location, resulting in multivariate spatial data. Let $\mathbf{x}(\mathbf{u})$ denote a p-variate random field with the cross-covariance function

$$\mathbf{C}_{\mathbf{x}}(\mathbf{u}_i, \mathbf{u}_j) = \mathrm{E}\big[(\mathbf{x}(\mathbf{u}_i) - \mathrm{E}[\mathbf{x}(\mathbf{u}_i)])(\mathbf{x}(\mathbf{u}_j) - \mathrm{E}[\mathbf{x}(\mathbf{u}_j)])^\top\big] .$$

Constructing flexible and valid multivariate covariance models is substantially more challenging than in the univariate case, as reviewed in Genton and Kleiber (2015). A widely used framework is the linear model of coregionalization (LMC) (Goulard and Voltz, 1992; Wackernagel, 2003), which represents the cross-covariance as a finite sum of coregionalization matrices multiplied by univariate spatial correlation functions. Multivariate extensions of the Matérn covariance have been proposed by Gneiting et al. (2010) and Apanasovich et al. (2012), though validity requires nontrivial parameter constraints.

Despite their popularity, these models rely on assumptions of global weak stationarity and isotropy, which are often unrealistic for large or heterogeneous spatial domains. Stationarity may be plausible locally, but it rarely holds globally. This has motivated extensive work on nonstationary spatial models, particularly in the univariate setting. Prominent approaches include spatial deformations (Sampson and Guttorp, 1992; Sampson, 2010), kernel convolution and closed-form nonstationary covariances (Paciorek and Schervish, 2006), spectral-domain methods (Fuentes, 2002), locally stationary estimation (Anderes and Stein, 2011), and SPDE-based (Stochastic Partial Differential Equation) models with spatially varying dependence parameters (Lindgren et al., 2011; Ingebrigtsen et al., 2014).

Developing nonstationary models for multivariate spatial data is even more demanding, as valid cross-covariances must simultaneously accommodate spatially varying marginal and cross dependence (Vu et al., 2022). Extensions of the LMC with spatially varying parameters (Gelfand et al., 2004) and nonstationary multivariate Matérn models (Kleiber and Nychka, 2012) have been proposed, as well as multivariate spatial deformation approaches (Vu et al., 2022). Nevertheless, these methods still rely on explicit specification of (possibly nonstationary) cross-covariance structures and become increasingly impractical as the dimension p grows.

Indeed, as p increases, the number of marginal and cross-covariance functions grows quadratically, leading to severe computational and identifiability issues. These challenges have fueled interest in *dimension reduction* techniques, which model multivariate spatial fields through a smaller number of latent components capturing the dominant dependence structure. Besides its role as a covariance construction, the LMC may be viewed as an implicit dimension reduction approach. More explicit methods include *spatial blind source separation* (SBSS) (Nordhausen et al., 2015; Bachoc et al., 2020; Muehlmann et al., 2024), which assumes that the observed field is a linear mixture of latent, second-order uncorrelated spatial sources with distinct dependence structures. Estimation typically relies on the joint diagonalization of the spatial covariance matrices across multiple lags, after which the latent components can be modeled individually. Extensions allowing for spatial nonstationarity, often referred to as *spatial nonstationary source separation* (SNSS) (Muehlmann et al., 2022; Sipilä et al., 2024), permit spatially varying variance or dependence structures of the latent sources.

Despite their advantages, existing source separation methods impose stationarity assumptions uniformly across all latent components, or require all components to be nonstationary in a similar manner. To the best of our knowledge, the possibility that the latent space decomposes into a stationary subspace and a nonstationary subspace has not yet been systematically explored in the spatial statistics literature. This stands in marked contrast to the time series context, where *stationary subspace analysis* (SSA) (see e.g. von Bünau et al., 2009; Hara et al., 2010; Flumian et al., 2024) is a well-established and actively researched framework for decomposing multivariate signals into stationary and nonstationary components. Such a decomposition is nevertheless appealing in spatial applications, where large-scale latent factors may induce nonstationarity, while smaller-scale components remain approximately stationary.

The goal of this paper is to develop *spatial stationary subspace analysis* (spSSA), a novel dimension reduction framework for multivariate spatial data that explicitly separates stationary and nonstationary latent subspaces. The proposed methodology allows for the detection of nonstationarities in the mean, variance, and spatial dependence structure. By isolating stationary components, spSSA enables their efficient analysis using standard univariate spatial models, while the nonstationary components—being of reduced dimension—can be modeled with substantially simplified nonstationary structures.

Unlike much of the existing stationary subspace analysis literature in time series, where stationary and nonstationary subspace dimensions are usually assumed to be known, we propose a novel estimator for the latent subspace dimension. This aspect is of independent interest and not only essential in practical spatial applications, but also readily transferable to the time series context.

The remainder of the paper is organized as follows. Section 2 introduces the necessary definitions and notation as well as presents the spSSA methods when the nonstationary dimension is known. Section 4 presents methods for finding said dimension using data augmentation. These methods are demonstrated in practice via simulation studies in Section 3 and Section 5, respectively, as well as using real data in Section 6. Sections A and B show additional simulation results.

2 Stationary subspace analysis

Let $\mathbf{x}(\mathbf{u})$ be an observable p-variate nonstationary random field, $\mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}^k$, where $\mathcal{U}$ denotes the spatial domain. Most commonly the dimension of the domain is either $k = 2$ or $k = 3$. The methods we present generalize for arbitrary k but, for our purposes, we fix $k = 2$.

Let us further assume that $\mathbf{x}$ can be decomposed into a stationary part and a nonstationary part as follows:

$$\mathbf{x}(\mathbf{u}) = \mathbf{A}\mathbf{z}(\mathbf{u}) + \boldsymbol{\mu} = [\mathbf{A_s}\ \mathbf{A_n}] \begin{pmatrix} \mathbf{s}(\mathbf{u}) \\ \mathbf{n}(\mathbf{u}) \end{pmatrix} + \boldsymbol{\mu} = \mathbf{A_s}\mathbf{s}(\mathbf{u}) + \mathbf{A_n}\mathbf{n}(\mathbf{u}) + \boldsymbol{\mu}, \tag{1}$$

where a p-variate latent random $\mathbf{z}(\mathbf{u})$ consists of a $(p-q)$-variate stationary random field $\mathbf{s}(\mathbf{u})$ and a q-variate nonstationary random field $\mathbf{n}(\mathbf{u})$, and $\boldsymbol{\mu}$ is a location vector. Without loss of generality, we can assume that $\boldsymbol{\mu} = \mathbf{0}$, and will do so henceforth. 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 $p \times (p-q)$ and $p \times q$ matrices, respectively. The aim of spatial stationary subspace analysis (spSSA) is to estimate an unmixing matrix $\mathbf{W}$ such that $\mathbf{W}\mathbf{x}$ is partitioned into stationary and nonstationary random fields.

We make the following assumptions on the latent random fields.

- (A1) $\mathrm{E}[\mathbf{s}(\mathbf{u})] = \mathbf{0}$, $\mathrm{Cov}(\mathbf{s}(\mathbf{u})) = \mathbf{I}$ for all $\mathbf{u}$, and $\mathrm{Cov}\left(\mathbf{s}(\mathbf{u}), \mathbf{s}(\mathbf{u}')\right)$ is finite and only a function of $\|\mathbf{u} - \mathbf{u}'\| = \|\mathbf{h}\|$.
- (A2) $\mathrm{E}[\mathbf{n}(\mathbf{u})] < \infty$, $\mathrm{Cov}(\mathbf{n}(\mathbf{u})) = \mathbf{D_u}$, where $\mathbf{D_u}$ is a diagonal matrix with positive diagonal elements, and $\mathrm{Cov}\left(\mathbf{n}(\mathbf{u}), \mathbf{n}(\mathbf{u}')\right)$ is finite.
- (A3) The random fields $\mathbf{s}(\mathbf{u})$ and $\mathbf{n}(\mathbf{u})$ are independent, i.e., $\mathrm{Cov}(\mathbf{s}(\mathbf{u}), \mathbf{n}(\mathbf{u}')) = \mathbf{0}$ for all $\mathbf{u}$ and $\mathbf{u}'$ in $\mathcal{U}$.
- (A4) The dimension q is the smallest value that allows division of $\mathbf{x}$ such that (A1)-(A3) hold.

We thus assume that the $(p-q)$ random fields in $\mathbf{s}(\mathbf{u})$ are second-order stationary with finite second order spatial dependence. The first assumption fixes the location and covariance matrix of the stationary part for convenience. The second assumption states that the q nonstationary components in $\mathbf{n}(\mathbf{u})$ have finite first and second moments and are uncorrelated at $\mathbf{u}$ but can have arbitrary location-dependent second order dependence when $\mathbf{u} \neq \mathbf{u}'$. The third assumption states that nonstationary and stationary components are independent. The last assumption ensures that $\mathbf{n}$ does not include stationary components.

Despite Assumptions (A1)–(A4), the model is not well-defined. The stationary components are only specified up to a multiplication by an orthogonal matrix, whereas the nonstationary components can be marginally rescaled, shifted and also rotated. Therefore, for convenience of presentation, we make the following additional assumption on $\mathbf{n}(\mathbf{u})$.

- (A5) $\sum_{\mathbf{u}\in\mathbf{U}_0} \mathrm{E}[\mathbf{n}(\mathbf{u})] = \mathbf{0}$ and $\sum_{\mathbf{u}\in\mathbf{U}_0} \mathrm{Cov}(\mathbf{n}(\mathbf{u})) = \mathbf{I}$, where $\mathbf{U}_0 = \{\mathbf{u}_1, \ldots, \mathbf{u}_{n_0}\} \subset \mathcal{U}$ is a set of n_0 reference locations.

Hence, based on a set of reference points, location and scale are fixed. Furthermore, for convenience we will assume that if a realization of $\mathbf{x}$ is observed at n spatial locations, then these locations will serve as reference points.

Throughout the rest of this article, we assume that the field $\mathbf{x}$ is defined as in Equation (1) and satisfies Assumptions (A1)–(A5). To estimate the unmixing matrix $\mathbf{W}$, we compute the first and the second order statistics based on the full domain $\mathcal{U}$ and then compare these statistics to the ones computed on partitions of the domain. To this end, let $\mathcal{U}_1, \ldots, \mathcal{U}_K$ be a disjoint partition of the domain $\mathcal{U}$, that is, $\mathcal{U}_i \cap \mathcal{U}_j = \varnothing$ for $1 \leq i \neq j \leq K$ and $\bigcup_{k\in K} \mathcal{U}_k = \mathcal{U}$. For the stationary signals we should observe the same statistic locally and globally. Conversely, for the nonstationary signals the global statistic should differ from the local statistics.

Suppose that the field $\mathbf{x}$ is observed at n locations $\mathbf{u}_1, \ldots, \mathbf{u}_n$. Define $\mathbf{U}_k = \mathbf{U}_0 \bigcap \mathcal{U}_k$ to be the set of locations inside the kth subdomain, and let $|\mathbf{U}_k|$ denote its cardinality, for $k = 1, \ldots, K$. Denote the local mean, local covariance, and local spatial covariance, respectively, for subdomain $\mathbf{U}_k$, by

$$\begin{aligned}
\mathbf{m}_{\mathbf{U}_k}(\mathbf{x}) &= \frac{1}{|\mathbf{U}_k|} \sum_{\mathbf{u}\in\mathbf{U}_k} \mathbf{x}(\mathbf{u}), \\
\mathrm{Cov}_{\mathbf{U}_k}(\mathbf{x}) &= \frac{1}{|\mathbf{U}_k|} \sum_{\mathbf{u}\in\mathbf{U}_k} \left(\mathbf{x}(\mathbf{u}) - \mathbf{m}_{\mathbf{U}_k}(\mathbf{x})\right) \left(\mathbf{x}(\mathbf{u}) - \mathbf{m}_{\mathbf{U}_k}(\mathbf{x})\right)^\top, \quad \text{and} \\
\mathrm{Lcov}_{\mathbf{U}_k,f}(\mathbf{x}) &= \frac{1}{|\mathbf{U}_k|} \sum_{\substack{\mathbf{u},\mathbf{u}'\in\mathbf{U}_k, \\ \mathbf{u}\neq\mathbf{u}'}} f(\mathbf{u}-\mathbf{u}') \left(\mathbf{x}(\mathbf{u}) - \mathbf{m}_{\mathbf{U}_k}(\mathbf{x})\right) \left(\mathbf{x}(\mathbf{u}') - \mathbf{m}_{\mathbf{U}_k}(\mathbf{x})\right)^\top,
\end{aligned}$$

where $f : \mathbb{R}^2 \to \mathbb{R}$ denotes a spatial kernel function. Options, as introduced in (Bachoc et al., 2020), include:

Ball kernel: $f_b(\mathbf{h}; r) = I(\|\mathbf{h}\| \leq r)$, where $r \geq 0$ and I is the indicator function.

Ring kernel: $f_r(\mathbf{h}; r_1, r_2) = I(r_1 < \|\mathbf{h}\| \leq r_2)$, where $r_1, r_2 \geq 0$ and $r_1 < r_2$.

Gaussian kernel: $f_g(\mathbf{h}; r) = \exp(-0.5(\Phi^{-1}(0.95)\|\mathbf{h}\|/r)^2)$, where $r > 0$ and $\Phi^{-1}(0.95)$ is the 95% quantile of the standard normal distribution.

As is the standard in blind source separation, the first step of the analysis is to whiten the data (Nordhausen and Ruiz-Gazen, 2022). Whitening reduces the problem to the estimation of an orthogonal transformation, since after standardization only an orthogonal matrix is required to recover the stationary and nonstationary subspaces. Accordingly, to separate $\mathbf{s}$ and $\mathbf{n}$, we work with the whitened random field defined by

$$\mathbf{x}^{st}(\mathbf{u}) = \mathrm{Cov}_{\mathbf{U}}(\mathbf{x})^{-1/2}(\mathbf{x}(\mathbf{u}) - \mathbf{m}_{\mathbf{U}}(\mathbf{x})).$$

2.1 Nonstationarity in mean

Consider first a spSSA method that aims at detecting stationary deviations in mean. Let $\mathbf{x}^{st}$ be the whitened field corresponding to $\mathbf{x}$ and let

$$\mathbf{M}_{\mathrm{mean}} = \sum_{k=1}^{K} \frac{|\mathbf{U}_k|}{|\mathbf{U}|} \mathbf{m}_{\mathbf{U}_k}(\mathbf{x}^{st}) \mathbf{m}_{\mathbf{U}_k}(\mathbf{x}^{st})^\top$$

be the covariance matrix of means computed on different domains weighted by their number of observations.

If the mean of a random field does not change in space, then its local means are approximately the same as the global mean, i.e., zero by Assumption (A1). The rows and columns, and hence also the diagonals of $\mathbf{M}_{\mathrm{mean}}$ corresponding to such fields are zero. Thus, the components that are stationary in mean, correspond to zero eigenvalues of $\mathbf{M}_{\mathrm{mean}}$. Conversely, the components that are nonstationary in mean correspond to non-zero eigenvalues. Therefore, write the eigenvalue decomposition of $\mathbf{M}_{\mathrm{mean}}$ as

$$\mathbf{M}_{\mathrm{mean}} = \mathbf{V}_{\mathrm{mean}} \mathbf{D}_{\mathrm{mean}} \mathbf{V}_{\mathrm{mean}}^\top,$$

where $\mathbf{D}_{\mathrm{mean}}$ is a $p \times p$ diagonal matrix of eigenvalues of $\mathbf{M}_{\mathrm{mean}}$, and $\mathbf{V}_{\mathrm{mean}}$ is the corresponding matrix of eigenvectors. Assume further that the eigenvalues and vectors are arranged so that $\mathbf{V}_{\mathrm{mean}} = (\mathbf{V}_{\mathrm{mean},s}\ \mathbf{V}_{\mathrm{mean},n})$ where $\mathbf{V}_{\mathrm{mean},n}$ is the $p \times q_{\mathrm{mean}}$ matrix containing the eigenvectors of non-zero eigenvalues as columns and the $p \times (p - q_{\mathrm{mean}})$ matrix $\mathbf{V}_{\mathrm{mean},s}$ includes the remaining columns. Then, columns of resulting unmixing matrices $\mathbf{W}_{\mathrm{mean},n} = \mathbf{V}_{\mathrm{mean},n}^\top \mathrm{Cov}_{\mathbf{U}}(\mathbf{x})^{-1/2}$ and $\mathbf{V}_{\mathrm{mean},s} = \mathbf{V}_{\mathrm{mean},s}^\top \mathrm{Cov}_{\mathbf{U}}(\mathbf{x})^{-1/2}$ generate the nonstationary and stationary subspaces, respectively. Naturally, $q_{\mathrm{mean}} \leq q$ with equality only if, for all nonstationary components, the means differ between at least some of the chosen partitions. We refer to this method as SPSSASIR.

2.2 Nonstationarity in variance

Consider again the whitened random field $\mathbf{x}^{st}$. Now

$$\mathbf{M}_{\mathrm{var}} = \sum_{k=1}^{K} \frac{|\mathbf{U}_k|}{|\mathbf{U}|} \left(\mathbf{I} - \mathrm{Cov}_{\mathbf{U}_k}(\mathbf{x}^{st})\right)^2,$$

where $\mathbf{A}^2 = \mathbf{A}\mathbf{A}^\top$, measures the deviation of the covariance matrix computed on partitions $\mathbf{U}_1, \ldots, \mathbf{U}_K$ from the global covariance matrix $\mathbf{I}_p$.

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

$$\mathbf{M}_{\mathrm{var}} = \mathbf{V}_{\mathrm{var}} \mathbf{D}_{\mathrm{var}} \mathbf{V}_{\mathrm{var}}^\top,$$

where $\mathbf{D}_{\mathrm{var}}$ is a $p \times p$ diagonal matrix of eigenvalues of $\mathbf{M}_{\mathrm{var}}$, and $\mathbf{V}_{\mathrm{var}} = (\mathbf{V}_{\mathrm{var},s}\ \mathbf{V}_{\mathrm{var},n})$ are the eigenvectors of $\mathbf{M}_{\mathrm{var}}$ arranged so that $\mathbf{V}_{\mathrm{var},n}$ is the $p \times q_{\mathrm{var}}$ matrix containing the eigenvectors corresponding to the non-zero eigenvalues as columns and $\mathbf{V}_{\mathrm{var},s}$ is the $p \times (p - q_{\mathrm{var}})$ matrix containing the remaining eigenvectors as columns. Again $q_{\mathrm{var}} \leq q$. The transformations to the two subspaces are accordingly given by the matrices $\mathbf{W}_{\mathrm{var},n} = \mathbf{V}_{\mathrm{var},n}^\top \mathrm{Cov}_{\mathbf{U}}(\mathbf{x}^{st})^{-1/2}$ and $\mathbf{W}_{\mathrm{var},s} = \mathbf{V}_{\mathrm{var},s}^\top \mathrm{Cov}_{\mathbf{U}}(\mathbf{x}^{st})^{-1/2}$. 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 partitions might also change. We refer to this method as SPSSASAVE.

2.3 Nonstationarity in spatial dependence

The two previous methods are able to detect components which are nonstationary with respect to the first and the second moments. To detect nonstationarities in spatial dependence, we need a statistic that measures the variability of spatial dependence across subdomains. Let such a statistic be defined (for whitened fields) as

$$\mathbf{M}_{\mathrm{cor},f} = \sum_{k=1}^{K} \frac{|\mathbf{U}_k|}{|\mathbf{U}|} \left(\mathrm{Lcov}_{\mathbf{U},f}(\mathbf{x}^{st}) - \mathrm{Lcov}_{\mathbf{U}_k,f}(\mathbf{x}^{st})\right)^2.$$

The matrix $\mathbf{M}_{\mathrm{cor},f}$ measures the deviation of the local covariance matrices computed on partitions $\mathbf{U}_1, \ldots, \mathbf{U}_K$ from the global local covariance matrix $\mathrm{Lcov}_{\mathbf{U},f}$. Using again the eigendecomposition

$$\mathbf{M}_{\mathrm{cor},f} = \mathbf{V}_{\mathrm{cor}}\mathbf{D}_{\mathrm{cor}}\mathbf{V}_{\mathrm{cor}}^{\top}, \tag{2}$$

and separating the eigenvectors of $\mathbf{M}_{\mathrm{cor},f}$ corresponding to non-zero and zero eigenvalues yields the $p \times q_{\mathrm{cor}}$ matrix $\mathbf{V}_{\mathrm{cor},n}$ and the $p \times (p - q_{\mathrm{cor}})$ matrix $\mathbf{V}_{\mathrm{cor},s}$, where q_{cor} is the number of non-zero eigenvalues with $q_{\mathrm{cor}} \leq q$. The unmixing matrix estimates are $\mathbf{W}_{\mathrm{cor},n} = \mathbf{V}_{\mathrm{cor},n}^{\top}\mathrm{Cov}_{\mathbf{U}}(\mathbf{x}^{st})^{-1/2}$ and $\mathbf{W}_{\mathrm{cor},s} = \mathbf{V}_{\mathrm{cor},s}^{\top}\mathrm{Cov}_{\mathbf{U}}(\mathbf{x}^{st})^{-1/2}$. The method is able to detect nonstationary components if at least some subdomains exhibit differences in their local covariance structure. As this method is for detecting changes in the spatial correlation structure we denote it as SPSSACOR.

Notice that as $\mathbf{M}_{\mathrm{cor},f}$ is computed using a single kernel f, the performance of SPSSACOR will naturally depend on the choice of f. However, there is no obvious rule for selecting f in practice in the absence of additional information about the underlying processes. Some visual and exploratory guidelines for choosing f in the context of SBSS are discussed in Piccolotto et al. (2022). An alternative approach is to employ multiple kernels $f_1, \ldots, f_L$ and to perform joint diagonalization of $\mathbf{M}_{\mathrm{cor},f_1}, \ldots, \mathbf{M}_{\mathrm{cor},f_L}$, rather than solely relying on the eigendecomposition in Equation (2).

Note that all the three preceding methods are formulated as an eigenvalue problem of a matrix computed from the whitened data $\mathbf{x}^{st}$. Equivalently, they can be formulated as generalized eigenvalue problems for the matrix pair $(\mathbf{M}\mathrm{Cov}_{\mathbf{U}}(\mathbf{x})^{-1/2}, \mathrm{Cov}_{\mathbf{U}}(\mathbf{x})^{1/2})$.

2.4 Combination of methods

Denote from now on $\mathbf{M}_{\mathrm{mean}} = \mathbf{M}_1$, $\mathbf{M}_{\mathrm{var}} = \mathbf{M}_2$ and $\mathbf{M}_{\mathrm{cor},f} = \mathbf{M}_3$. The methods of Sections Sections 2.1 to 2.3 can be combined to detect all three types of nonstationarities by solving a joint diagonalization problem. Since exact joint diagonalization of more than two matrices is usually impossible, approximate joint diagonalization is required. For the three matrices $\mathbf{M}_i$, $i = 1, 2, 3$, computed with the whitened random field, this amounts to finding an orthogonal $p \times p$ matrix $\mathbf{V}_{\mathrm{comb}}$ that minimizes $\sum_{i=1}^{3} \left\|\mathrm{off}(\mathbf{V}_{\mathrm{comb}}^{\top}\mathbf{M}_i\mathbf{V}_{\mathrm{comb}})\right\|^2$, or, equivalently, since $\mathbf{V}_{\mathrm{comb}}$ is orthogonal, maximizes

$$\sum_{i=1}^{3} \left\|\mathrm{diag}(\mathbf{V}_{\mathrm{comb}}^{\top}\mathbf{M}_i\mathbf{V}_{\mathrm{comb}})\right\|^2. \tag{3}$$

Here $\|\mathbf{A}\|$ is the matrix (Frobenius) norm, $\mathrm{diag}(\mathbf{A})$ is a $p \times p$ diagonal matrix with the diagonal elements as in $\mathbf{A}$ and $\mathrm{off}(\mathbf{A}) = \mathbf{A} - \mathrm{diag}(\mathbf{A})$. Notice that the principle of the estimation procedure remains the same even if additional matrices $\mathbf{M}_{\mathrm{cor},f_i}$ with various kernels f_i are included in the objective function. Similarly, it is possible to only consider a subset of the matrices $\mathbf{M}_i$.

Several algorithms for approximate joint diagonalization in Equation (3) exist in the literature. The most popular one based on Givens rotations is proposed in Clarkson (1988) and is available for example in the R package JADE (Miettinen et al., 2017).

Based on $\mathbf{V}_{\mathrm{comb}}$, one can compute $\mathbf{D}_i = \mathrm{diag}(\mathbf{V}_{\mathrm{comb}}\mathbf{M}_i\mathbf{V}_{\mathrm{comb}}^{\top}) = \mathrm{diag}(d_{i,1}, \ldots, d_{i,p})$ and collect all those columns j of $\mathbf{V}_{\mathrm{comb}}$, where $\sum_i |d_{i,j}| \neq 0$, to a $p \times q_{\mathrm{comb}}$ matrix $\mathbf{V}_{\mathrm{comb},n}$. The rest of the columns are then collected to a $p \times (p - q_{\mathrm{comb}})$ matrix $\mathbf{V}_{\mathrm{comb},s}$. Here the individual value $d_{i,j}$ indicates whether the jth component is nonstationary with respect to $\mathbf{M}_i$. The final transformation matrices for nonstationary and stationary components are then $\mathbf{W}_{\mathrm{comb},n} = \mathbf{V}_{\mathrm{comb},n}^{\top}\mathrm{Cov}_{\mathbf{U}}(\mathbf{x}^{st})^{-1/2}$ and $\mathbf{W}_{\mathrm{comb},s} = \mathbf{V}_{\mathrm{comb},s}^{\top}\mathrm{Cov}_{\mathbf{U}}(\mathbf{x}^{st})^{-1/2}$. We refer to this method as SPSSACOMB.

Remark 1. Alternatively, one may consider forming a linear combination of the matrices $\mathbf{M}_1, \mathbf{M}_2$, and $\mathbf{M}_3$ and performing an eigenvalue decomposition of the resulting matrix. This approach, however, introduces nontrivial scaling issues. For example, in a time series context, Hara et al. (2010) consider a weighted linear combination of $\mathbf{M}_1$ and $\mathbf{M}_2$. Such weighting becomes necessary when one of the matrices has a substantially larger norm than the others; without appropriate scaling, the combined matrix would yield results that differ only marginally from those obtained using the dominant matrix alone. A seemingly straightforward remedy is to scale each matrix by an appropriate matrix norm.

However, this can be problematic: if a particular type of nonstationarity is absent, the corresponding scatter matrix contains only noise, and dividing by its nonzero, but very small norm, may strongly amplify the noise. For this reason, we do not consider this approach in the simulation study. We note that similar scaling issues may also arise in the context of approximate joint diagonalization.

Remark 2. As noted by Flumian et al. (2024) in the time series setting, the SSA approaches considered here may also be viewed through the lens of supervised dimension reduction (SDR). From this perspective, an artificial response variable is given by the partition membership, and the methods aim to identify components that best discriminate between partitions, that is, the nonstationary components. This interpretation also explains the terminology SIR and SAVE, which is motivated by the corresponding SDR methods sliced inverse regression (Li, 1991) and sliced average variance estimation (Cook, 2000).

2.5 Practical considerations

The choice of the spatial partition plays a crucial role in the performance of the spSSA procedures. An unsuitable partition may cause the methods to indicate the absence of nonstationarity, even when pronounced local differences are present. This is closely related to the *modifiable areal unit problem* (MAUP), see e.g., Openshaw (1984); Fotheringham and Wong (1991). An illustrative example is shown in Figure 1, which displays two subfigures with the same univariate random field with 450 data points. The range of values is split into seven groups, and the value of the random field at a given location is represented by the label of the group. The ranges and labels are given in the legend. Globally, the data has mean 0 and variance 1. While the random field is stationary in variance and covariance, its mean varies locally across the domain. The spatially varying mean is generated on a 3×3 grid, and using a corresponding 3×3 partition allows SPSSASIR to successfully detect this nonstationary structure. In contrast, when a 2×2 partition is used, the regional means do not differ, and the nonstationarity in the mean remains undetected.

A similar effect is illustrated for variance in Figure 2, where this time the random field is stationary in mean but nonstationary in variance. Analogous issues arise for nonstationarity in the spatial dependence structure. In general, partitions should be chosen to maximize differences between regions, a task for which subject-matter knowledge can be highly beneficial.

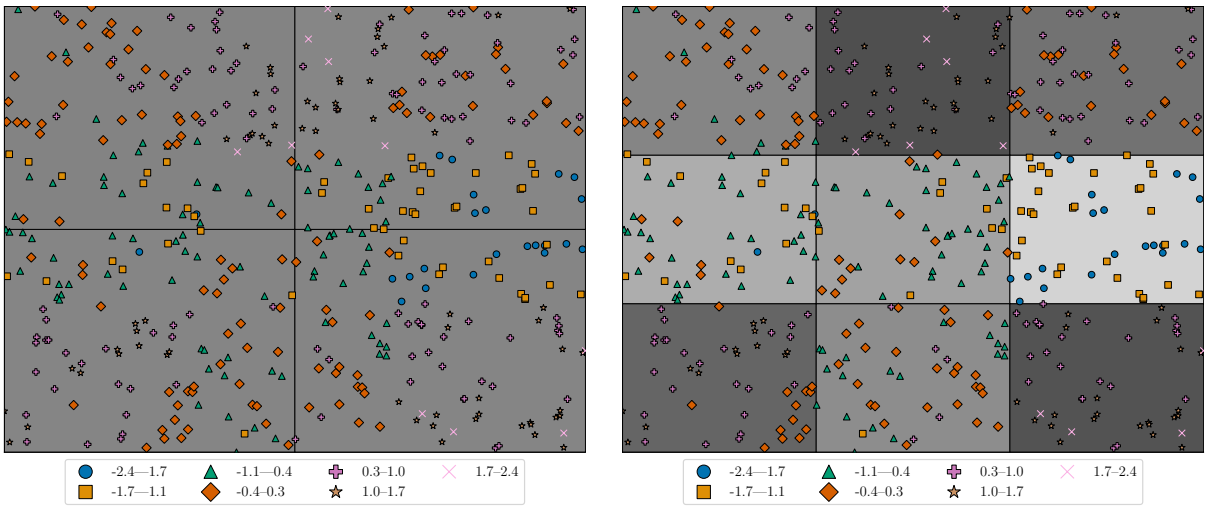

(a) 2-by-2 partition

(b) 3-by-3 partition

Figure 1: Example of the importance of choosing the right partition for a spatial signal that is nonstationary in mean. The labels represent ranges of observed values of the spatial signal. The background color represents the mean of a box. The darker the color the higher the mean.

As noted in Remark 1, the relative scaling of the $\mathbf{M}$ matrices can substantially affect the performance of SPSSACOMB. To fully exploit approximate joint diagonalization, it is important that the norms of these matrices are of comparable magnitude. If one matrix dominates in norm, the resulting decomposition would essentially recover the eigenstructure of the dominant matrix, with only a negligible contribution from the remaining matrices.

After whitening the signals, the matrices $\mathbf{M}_{\text{mean}}$ and $\mathbf{M}_{\text{var}}$ are expected to be of comparable norm. This is not necessarily the case for $\mathbf{M}_{\text{cor}}$. For purely stationary signals, local means and variances should closely match their

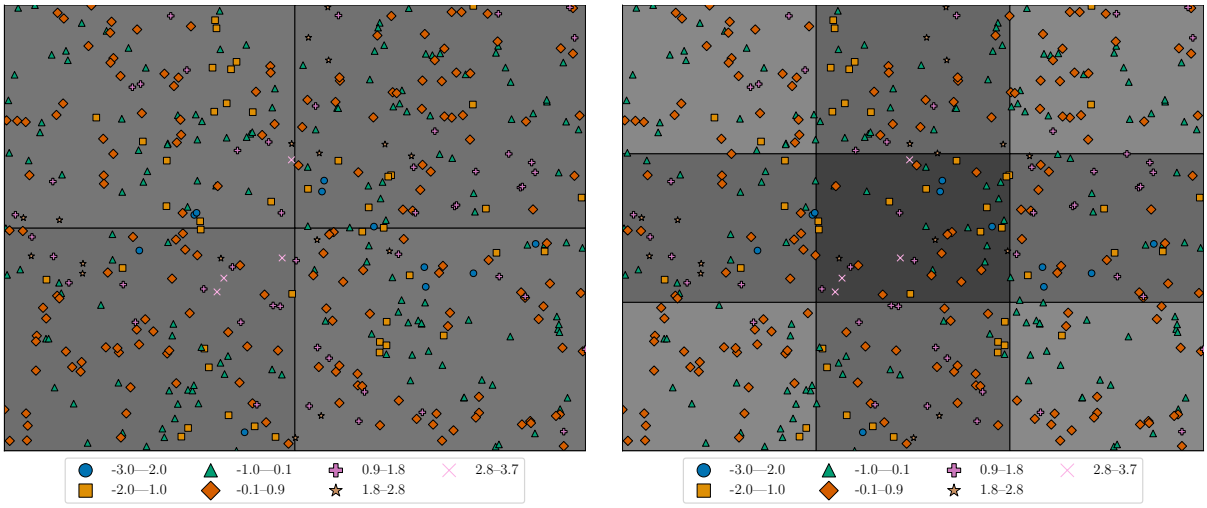

(a) 2-by-2 partition

(b) 3-by-3 partition

Figure 2: Example of the importance of choosing the right partition for a spatial signal that is nonstationary in variance. The labels represent ranges of observed values of the spatial signal. The background color represents the variance of a box. The darker the color the higher the variance.

global counterparts, implying that the norms of $\mathbf{M}_{\text{mean}}$ and $\mathbf{M}_{\text{var}}$ remain small. In contrast, local and global spatial covariances can differ significantly, even under stationarity, due to variations in the spatial configuration and density of measurement locations within partitions. This effect becomes more pronounced depending on the partition, and the norm of $\mathbf{M}_{\text{cor},f}$ can potentially grow alongside the number of partitions.

This motivates the introduction of scaling. By appropriately dampening the influence of densely sampled regions, scaling reduces exaggerated local–global differences in spatial dependence and keeps the norm of $\mathbf{M}_{\text{cor},f}$ on a scale comparable to that of $\mathbf{M}_{\text{mean}}$ and $\mathbf{M}_{\text{var}}$.

Let us define a scaled version of local covariance $\text{Lcov}^s_{\mathbf{U}_k,f}$ by

$$\text{Lcov}^s_{\mathbf{U}_k,f}(\mathbf{x}) = \frac{1}{|\mathbf{U}_k|} \sum_{\mathbf{u}\in\mathbf{U}_k} \frac{1}{F_{\mathbf{U}_k}(\mathbf{u};f)} \sum_{\mathbf{u}'\in\mathbf{U}_k\setminus\mathbf{u}} f(\mathbf{u}-\mathbf{u}')\left(\mathbf{x}(\mathbf{u})-\mathbf{m}_{\mathbf{U}_k}(\mathbf{x})\right)\left(\mathbf{x}(\mathbf{u}')-\mathbf{m}_{\mathbf{U}_k}(\mathbf{x})\right)^\top,$$

where

$$F_{\mathbf{U}_k}(\mathbf{u};f) = \sum_{\mathbf{u}'\in\mathbf{U}_k\setminus\mathbf{u}} f(\mathbf{u}-\mathbf{u}').$$

Analogously, define $\mathbf{M}^s_{\text{cor},f}$ by

$$\mathbf{M}^s_{\text{cor},f} = \sum_{k=1}^{K} \frac{|\mathbf{U}_k|}{|\mathbf{U}|} \left(\text{Lcov}^s_{\mathbf{U},f}(\mathbf{x}^{st}) - \text{Lcov}^s_{\mathbf{U}_k,f}(\mathbf{x}^{st})\right)^2.$$

The scaling factor accounts for the number of nearby locations contributing to the local covariance at a location $\mathbf{u}$. For the ball kernel, the scaling factor $F_{\mathbf{U}_k}(\mathbf{u}, f_b(\,\cdot\,;r))$ is simply the number of measurement locations contained in $B(\mathbf{u},r)\cap\mathbf{U}_k$. The scaled local covariance retains the same structural properties as the unscaled version, since the scaling factor can be absorbed into the kernel function.

The use of a scaled local covariance is not strictly required for any of the methodologies introduced in this paper. However, in the context of approximate joint diagonalization, it resolves the scaling and partitioning issues discussed above. In particular, the scaled local covariance allows $\mathbf{M}_{\text{comb}}$ to capture local covariance information without overwhelming the contributions of $\mathbf{M}_{\text{mean}}$ and $\mathbf{M}_{\text{var}}$. For this reason, without loss of generality, we employ the scaled local covariance estimator in the subsequent sections, while continuing to refer to the corresponding method as SPSSACOR. Note that there are also other scaled variants of the local covariance, as for example discussed in Muehlmann et al. (2024). A comparison between the scaled and unscaled SPSSACOR, as well as an example of the scaling issue, can bee seen in Section A.

3 Simulations I

In this section, we assess the performance of the proposed methods under the assumption that the subspace dimension q is known. For a given unmixing matrix $\mathbf{W} = [\mathbf{W}_s\ \mathbf{W}_n]$ and its estimate $\widetilde{\mathbf{W}} = \left[\widetilde{\mathbf{W}}_s\ \widetilde{\mathbf{W}}_n\right]$, we compute their respective projectors

$$\mathbf{P}_s = \mathbf{W}_s \left(\mathbf{W}_s^\top \mathbf{W}_s\right)^{-1} \mathbf{W}_s^\top, \qquad \mathbf{P}_n = \mathbf{W}_n \left(\mathbf{W}_n^\top \mathbf{W}_n\right)^{-1} \mathbf{W}_n^\top$$

and

$$\widetilde{\mathbf{P}}_s = \widetilde{\mathbf{W}}_s \left(\widetilde{\mathbf{W}}_s^\top \widetilde{\mathbf{W}}_s\right)^{-1} \widetilde{\mathbf{W}}_s^\top, \qquad \widetilde{\mathbf{P}}_n = \widetilde{\mathbf{W}}_n \left(\widetilde{\mathbf{W}}_n^\top \widetilde{\mathbf{W}}_n\right)^{-1} \widetilde{\mathbf{W}}_n^\top.$$

We choose, as a performance metric, the squared distance between the true and estimated projectors. That is, we denote by $\mathbf{s}_{\text{perf}}$ the stationary performance and $\mathbf{n}_{\text{perf}}$ the nonstationary performance, which are given by

$$\mathbf{s}_{\text{perf}} = \frac{1}{2}\left\|\mathbf{P}_s - \widetilde{\mathbf{P}}_s\right\|^2 \text{ and } \mathbf{n}_{\text{perf}} = \frac{1}{2}\left\|\mathbf{P}_n - \widetilde{\mathbf{P}}_n\right\|^2$$

For properties of the performance metric, see Liski et al. (2016). A key result is that the performance metric is upper bounded by the dimension of the subspace we are projecting to, that is, $\mathbf{s}_{\text{perf}} \in [0, \dim(\mathbf{s}(\mathbf{u}))]$ and $\mathbf{n}_{\text{perf}} \in [0, \dim(\mathbf{n}(\mathbf{u}))]$. Ideally, we want these metrics to converge to zero as we increase the number of measurements.

All of the following simulations are done with Python, and the source code is available at `github.com/perttusaarela/sp_ssa/releases/tag/v1.0`. In all settings, the observed random field $\mathbf{x}(\mathbf{u})$ is 8-dimensional with the stationary subspace $\mathbf{s}(\mathbf{u})$ begin 5-dimensional and the nonstationary subspace $\mathbf{n}(\mathbf{u})$ having dimension 3. The observed locations are drawn uniformly at random from the domain $\mathcal{U} = [0, \ell]^2$, that is, a square area whose side length is $\ell \in \{20, 30, 40, 50, 60, 70\}$. The number of locations n is the side length squared. In this way, we keep approximately the same density as we increase the number of measurements. The latent components are mixed by a random 8×8 orthogonal matrix $\mathbf{A}$. This is done without loss of generality, since the data is standardized.

For the spatial dependencies between these measurement locations, we use the Matérn covariance matrix defined by the function

$$C(\mathbf{h}; \nu, \phi) = \frac{1}{2^{\nu-1}\Gamma(\|\mathbf{h}\|)} \left(\frac{\|\mathbf{h}\|}{\phi}\right)^{\nu} K_\nu \left(\frac{\|\mathbf{h}\|}{\phi}\right),$$

where $\|\cdot\|$ is the Euclidean norm, K_ν is a modified Bessel function of the second-kind, ν is the shape parameter and ϕ is the range parameter. The vector $\mathbf{h}$ is the difference between two measurement points.

Components $\mathbf{s}_i(\mathbf{u})$, $i \in \{1, 2, 3, 4, 5\}$, of the stationary signal $\mathbf{s}(\mathbf{u})$ are independent zero mean Gaussian processes, where the covariance structure is determined by the Matérn covariance function

$$\text{Cov}(\mathbf{s}_i(\mathbf{u}), \mathbf{s}_i(\mathbf{u}')) = C(\mathbf{u} - \mathbf{u}'; 0.5, 1.0),$$

which guarantees that Assumption (A1) is satisfied. For the signals $\mathbf{n}(\mathbf{u})$, we consider three types of nonstationarities. In the first three settings, only one type of nonstationarity is present. Settings 1, 2, and 3 have nonstationarity in mean, variance, and spatial covariance, respectively. Setting 4 will combine all of these nonstationarities.

All the methods above rely on a partition $\mathcal{U}_1, \ldots, \mathcal{U}_K$ of the domain. For the simulations, we use a simple grid partition. For each of the settings, we consider three partitions: 2-by-2, 3-by-3, and 4-by-4 (so $K = 4, 9, 16$ respectively). In subsequent plots a k-by-ℓ partition is denoted by (k, ℓ). We compute the averages of performance measures from 2000 trials to smooth out single run variability.

Finally, along with the introduced methods, we also plot the performance of a random baseline. This is computed by generating a random $n \times n$ orthogonal matrix $\mathbf{U}$ and setting the unmixing matrix to $\mathbf{W} = \mathbf{U}^\top \text{Cov}_{\mathbf{U}}(\mathbf{x}^{st})^{-1/2}$. The idea of the baseline is to demonstrate the efficacy of our methods compared to random guessing.

Setting 1 The stationary signal $\mathbf{s}(\mathbf{u})$ is defined as above. The nonstationary signal $\mathbf{n}(\mathbf{u})$ consists of three components of the form $\mathbf{n}_i(\mathbf{u}) = y_i(\mathbf{u}) + \mu_i(\mathbf{u})$, where $y_i(\mathbf{u})$ is a stationary signal generated in the same way as the components of $\mathbf{s}(\mathbf{u})$, and $\mu_i(\mathbf{u})$ is a spatially varying mean function. The function μ_i is a piecewise constant function whose values are defined by the partition $A^{(i)}$. The mean functions are

$$\mu_1(\mathbf{u}) = \begin{cases} 1.5 & \mathbf{u} \in A_1^{(1)} \text{ or } A_4^{(1)}, \\ -1.5 & \mathbf{u} \in A_2^{(1)} \text{ or } A_3^{(1)}, \end{cases} \mu_2(\mathbf{u}) = \begin{cases} 1 & \mathbf{u} \in A_1^{(2)} \text{ or } A_5^{(1)}, \\ -0.5 & \mathbf{u} \in A_2^{(2)} \text{ or } A_6^{(1)}, \\ 2.0 & \mathbf{u} \in A_3^{(2)} \text{ or } A_4^{(1)}, \end{cases} \mu_3(\mathbf{u}) = \begin{cases} -1.5 & \mathbf{u} \in A_1^{(3)} \text{ or } A_6^{(1)}, \\ -0.5 & \mathbf{u} \in A_2^{(3)} \text{ or } A_7^{(1)}, \\ 0.5 & \mathbf{u} \in A_3^{(3)} \text{ or } A_8^{(1)}, \\ 1.5 & \mathbf{u} \in A_4^{(3)} \text{ or } A_5^{(1)}. \end{cases}$$

Figure 3 shows the partitions $A^{(i)}$ which are reused for all settings. More concretely, the defining partition of a nonstationary signal $\mathbf{n}_i$ is a $(i+1)$-by-$(i+1)$ grid partition. To accentuate the results, we assign different parameters to neighboring parts of the partition. To this end, we reuse parameters for multiple parts. Note that these are the same partitions that we also use for the estimations.

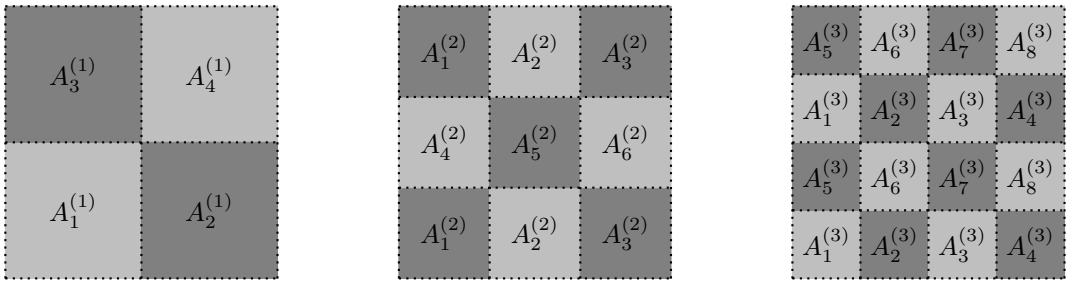

(a) Partition of parameters for signals n_1 (b) Partition of parameters for signals n_2 (c) Partition of parameters for signals n_3

Figure 3: Partitions for parameters used for generating the nonstationary signal $\mathbf{n}(\mathbf{u})$ in all settings.

Figure 4 shows that all the presented methods find nonstationarities in the mean. Based on how the signals are defined, SPSSASIR struggles slightly on the 2-by-2 partition but outperforms other methods when granularity is increased.

Setting 2 The stationary part $\mathbf{s}(\mathbf{u})$ is as defined above. The nonstationary part $\mathbf{n}(\mathbf{u})$ consists of three signals $\mathbf{n}_i$, where each signal is generated by scaling a stationary signal y with a regionally varying standard deviation. In other words, $\mathbf{n}_i(\mathbf{u}) = \sigma_i(\mathbf{u})y_i(\mathbf{u})$, where y_i is generated as in Setting 1 and the scaling functions are defined as

$$\sigma_1^2(\mathbf{u}) = \begin{cases} 0.4 & \mathbf{u} \in A_1^{(1)} \text{ or } A_4^{(1)}, \\ 1.4 & \mathbf{u} \in A_2^{(1)} \text{ or } A_3^{(1)}, \end{cases} \quad \sigma_2^2(\mathbf{u}) = \begin{cases} 3.0 & \mathbf{u} \in A_1^{(2)} \text{ or } A_5^{(1)}, \\ 0.5 & \mathbf{u} \in A_2^{(2)} \text{ or } A_6^{(1)}, \\ 1.5 & \mathbf{u} \in A_3^{(2)} \text{ or } A_4^{(1)}, \end{cases} \quad \sigma_3^2(\mathbf{u}) = \begin{cases} 0.4 & \mathbf{u} \in A_1^{(3)} \text{ or } A_6^{(1)}, \\ 0.8 & \mathbf{u} \in A_2^{(3)} \text{ or } A_7^{(1)}, \\ 1.5 & \mathbf{u} \in A_3^{(3)} \text{ or } A_8^{(1)}, \\ 1.2 & \mathbf{u} \in A_4^{(3)} \text{ or } A_5^{(1)}. \end{cases}$$

The regions for each signal are the same as the corresponding signal in Setting 1 and can be seen in Figure 3. This method for generating the data is identical to using the following covariance matrix for $\mathbf{n}_k(\mathbf{u})$:

$$\text{Cov}(\mathbf{n}_k(\mathbf{u}_i), \mathbf{n}_k(\mathbf{u}_j)) = \sigma_k(\mathbf{u}_i)\sigma_k(\mathbf{u}_j)C(\mathbf{u}_i - \mathbf{u}_j; 0.5, 1.0).$$

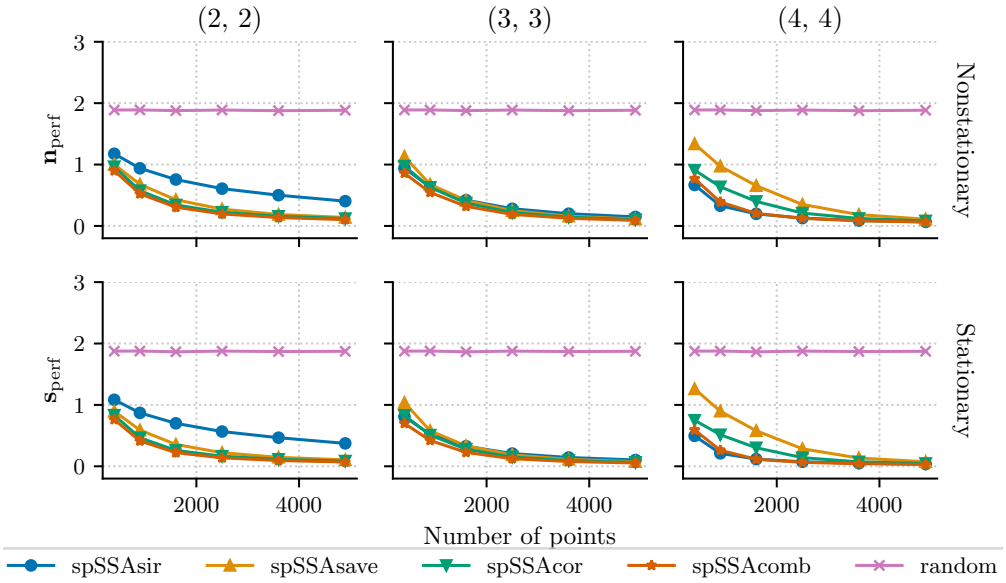

Figure 4: Setting 1: nonstationarity in the mean.

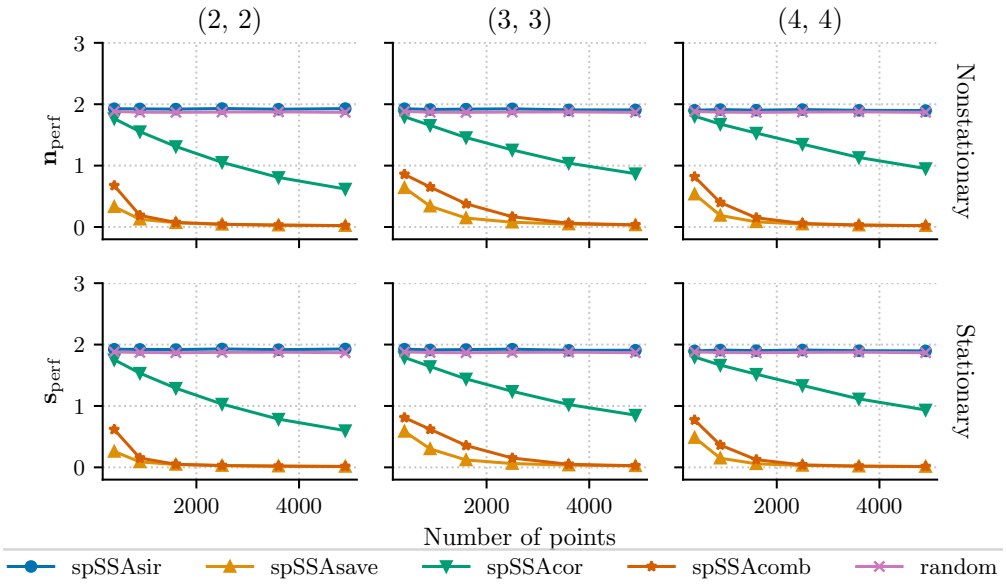

Figure 5: Setting 2: nonstationarity in variance

Note that this also causes some nonstationarity in the spatial covariance.

Figure 5 shows that SPSSASAVE finds all the signals even with few measurements. Since the mean is zero both globally and locally, SPSSASIR performs as bad as the baseline. With enough data points, SPSSACOR seems to be able to find some nonstationary signals. This is to be expected based on how we actually generate the data.

Setting 3 The stationary signals are again generated as above. For the nonstationary signals, each signal is generated independently. For each signal $\mathbf{n}_i(\mathbf{u})$, data are generated independently across partitions using Matérn covariances with partition-specific parameters (ν, ϕ) (see Table 1).

Table 1: Parameter values (ν, ϕ) for each block in the three partitions.

Partition	$A_1^{(\ell)}$	$A_2^{(\ell)}$	$A_3^{(\ell)}$	$A_4^{(\ell)}$	$A_5^{(\ell)}$	$A_6^{(\ell)}$	$A_7^{(\ell)}$	$A_8^{(\ell)}$
$\ell = 1$	(0.3, 0.5)	(1.5, 1.3)	(1.0, 2.0)	(0.5, 2.0)	–	–	–	–
$\ell = 2$	(1.0, 1.5)	(0.5, 0.8)	(2.0, 1.7)	(0.5, 2.0)	(1.0, 2.0)	(0.5, 2.0)	–	–
$\ell = 3$	(1.6, 1.6)	(0.3, 0.3)	(2.5, 3.0)	(0.8, 3.0)	(0.5, 1.8)	(1.0, 3.0)	(0.5, 1.2)	(0.3, 2.5)

We use the ball kernel $f_b(\mathbf{h}; r)$ with radius $r = 3.4$ for all simulations. The radius is chosen such that we have sufficient number of observations to estimate the covariance matrices. Assuming uniform density for the points, there should be approximately πr^2 points in a ball with radius r. Since we are estimating an 8×8 covariance matrix, we have 36 parameters to estimate. Setting $r = 3.4$ gives $\pi r^2 \approx 36.3$.

Figure 6 shows that SPSSASIR finds nothing, while SPSSASAVE finds some signals with enough data points. The best method is, as expected, SPSSACOR followed by SPSSACOMB. This also highlights the importance of the partition: even SPSSACOR does not fully converge to zero when the partition is too sparse.

Setting 4 The same stationary signals are used as above. The nonstationary signals, $\mathbf{n}_1$, $\mathbf{n}_2$, and $\mathbf{n}_3$ correspond to those of Settings 1, 2, and 3, respectively. We again use the ball kernel with radius $r = 3.4$. Figure 7 shows that the results follow our expectations based on the previous settings: SPSSASIR finds one nonstationary signal, SPSSASAVE and SPSSACOR find two, and SPSSACOMB converges to zero. For this set of parameters, each individual method performs the best on the 3-by-3 partition, but SPSSACOMB converges the fastest on the 4-by-4 partition. As long as there are enough points per part, increasing granularity increases performance.

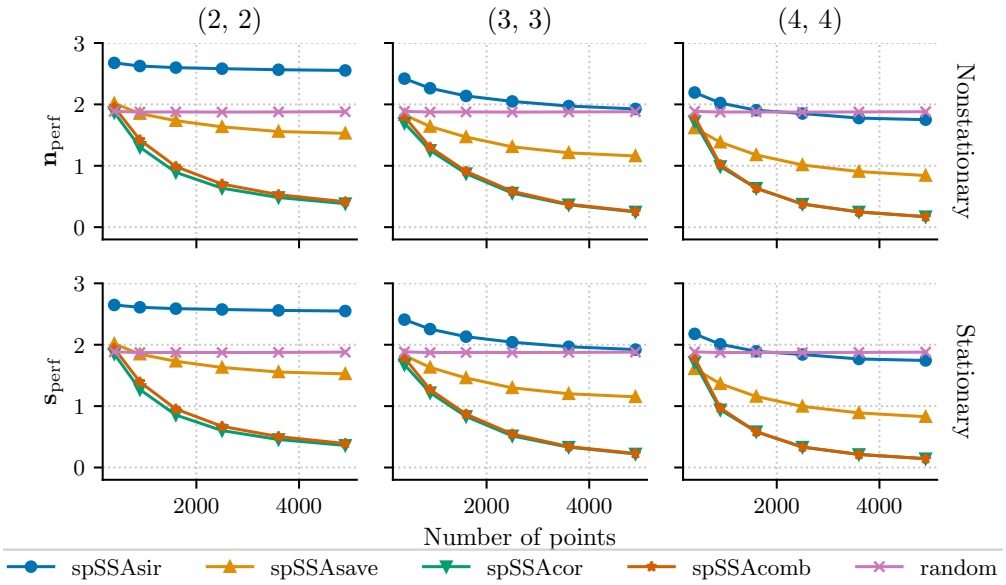

Figure 6: Setting 3: nonstationarity in spatial covariance.

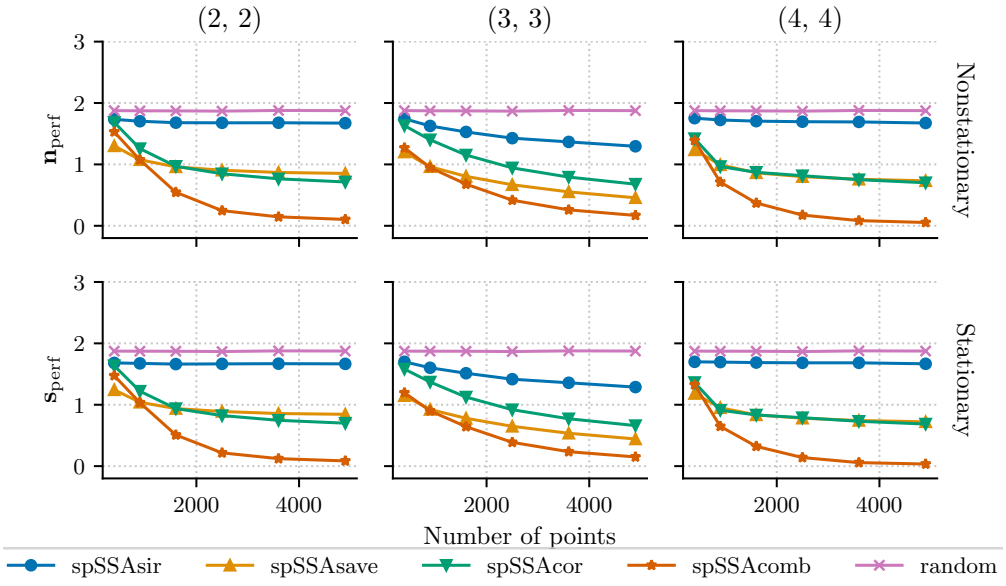

Figure 7: Setting 4: all three types of nonstationarity are present.

We have established that the proposed methods discriminate the two subspaces successfully under the unrealistic assumption that the dimensions of the subspaces are known. In the following, we suggest a procedure to estimate q.

4 Estimating the nonstationary subspace dimension with augmentation

So far, we have assumed that the subspace dimensions are known. However, in practice this assumption is clearly unrealistic. It is worth noting that even in the time series literature there is no generally established methodology for estimating these dimensions.

In the following, we introduce an estimator for the nonstationary subspace dimension. The estimator can be employed across all the spSSA methods discussed in Section 2. This is possible because the they share a common underlying structure:

1. Whiten the data to obtain $\mathbf{x}^{st}$.
2. Compute a matrix $\mathbf{M}(\mathbf{x}^{st})$, which is expected to have q non-zero eigenvalues and $p-q$ zero eigenvalues. Since the signs of the eigenvalues do not matter for our purposes, we consider their absolute values and assume the ordering
$$d_1 \geq \cdots \geq d_q > 0 = \cdots = 0.$$
3. Decompose the matrix as
$$\mathbf{M}(\mathbf{x}^{st}) = \mathbf{V}\mathbf{D}\mathbf{V}^\top,$$
where $\mathbf{D}$ is the diagonal matrix of eigenvalues and $\mathbf{V}$ contains the corresponding orthonormal eigenvectors.

In the case of Section 2.4, we get multiple $\mathbf{M}$ matrices, so the framework does not apply directly. However, we do get a decomposition in terms of pseudo-eigenvalues $\mathbf{D}$, and the pseudo-eigenvector matrix $\mathbf{V}$ which are the necessary components for the following.

Based on the structure of $\mathbf{M}$, we estimate the nonstationary subspace dimension q with the so-called augmentation approach which was originally proposed by Luo and Li (2021) in the context of dimension reduction for i.i.d. vector-valued data. This framework has since been extended to i.i.d. matrix- and tensor-valued observations (Radojičić et al., 2021, 2025), as well as to time series data (Radojičić and Nordhausen, 2024). Nevertheless, to the best of our knowledge, no corresponding extension to spatial data or to the setting of stationary subspace analysis has been developed.

Augmentation procedure. Let $r \geq 1$ be fixed. Let $\boldsymbol{\varepsilon}(\mathbf{u}) \in \mathbb{R}^r$ denote a white noise random field, independent of $\mathbf{x}(\mathbf{u})$. This is to say that $\boldsymbol{\varepsilon}(\mathbf{u})$ is a stationary random field satisfying Assumption (A1). The white noise random field is used to construct the *augmented standardized random field*

$$\mathbf{x}^{st}(\mathbf{u})^* = \begin{pmatrix} \mathbf{x}^{st}(\mathbf{u}) \\ \boldsymbol{\varepsilon}(\mathbf{u}) \end{pmatrix}.$$

Under the assumed model, the associated matrix becomes

$$\mathbf{M}(\mathbf{x}^{st^*}) = \mathbf{V}^* \begin{pmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_{\boldsymbol{\varepsilon}} \end{pmatrix} \mathbf{V}^{*\top} = \mathbf{V}^* \begin{pmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \mathbf{V}^{*\top},$$

where $\mathbf{V}^*$ is the diagonalizer of $\mathbf{M}(\mathbf{x}^{st^*})$ and $\mathbf{D}_{\boldsymbol{\varepsilon}}$ is a diagonal matrix corresponding to the eigenvalues associated with the white noise. Since the white noise is stationary, we know that the associated eigenvalues are zero. Furthermore, this tells us that $\operatorname{rank}(\mathbf{M}(\mathbf{x}^{st^*})) = \operatorname{rank}(\mathbf{M}(\mathbf{x}^{st})) = q$, i.e., the augmentation does not affect the rank of $\mathbf{M}$.

Actually, we can say even more about the structure of the augmented matrices. For each of the scatter matrices $\mathbf{M}_{\text{mean}}$, $\mathbf{M}_{\text{var}}$, $\mathbf{M}_{\text{cor},f}$, it holds that

$$\mathbf{M}(\mathbf{x}^{st^*}) = \begin{pmatrix} \mathbf{M}(\mathbf{x}^{st}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

This follows from the properties of scatter matrices and the assumptions on the present signals. For example, in the case of $\mathbf{M}_{\text{var}}$, we can examine each summand separately. For each $\mathbf{U}_k$, the covariance matrix $\operatorname{Cov}_{\mathbf{U}_k}(\mathbf{x}^{st^*})$ is given by

$$\mathrm{E}\left[\begin{pmatrix} \mathbf{x}^{st}(\mathbf{u}) - \mathrm{E}[\mathbf{x}^{st}(\mathbf{u})] \\ \boldsymbol{\varepsilon}(\mathbf{u}) - \mathrm{E}[\boldsymbol{\varepsilon}(\mathbf{u})] \end{pmatrix} \begin{pmatrix} \mathbf{x}^{st}(\mathbf{u})^\top - \mathrm{E}[\mathbf{x}^{st}(\mathbf{u})^\top] \\ \boldsymbol{\varepsilon}(\mathbf{u})^\top - \mathrm{E}[\boldsymbol{\varepsilon}(\mathbf{u})^\top] \end{pmatrix}^\top\right] = \mathrm{E}\left[\begin{pmatrix} \mathbf{x}^{st}(\mathbf{u})\mathbf{x}^{st}(\mathbf{u})^\top & \mathbf{x}^{st}(\mathbf{u})\boldsymbol{\varepsilon}(\mathbf{u})^\top \\ \boldsymbol{\varepsilon}(\mathbf{u})\mathbf{x}^{st}(\mathbf{u})^\top & \boldsymbol{\varepsilon}(\mathbf{u})\boldsymbol{\varepsilon}(\mathbf{u})^\top \end{pmatrix}\right]. \tag{4}$$

Looking at the matrix on the right, we make the following observations: the top left block is simply $\mathrm{Cov}(\mathbf{x}^{st})$. Next, since $\boldsymbol{\varepsilon}(\mathbf{u})$ is independent of $\mathbf{x}^{st}(\mathbf{u})$,

$$\mathrm{E}[\mathbf{x}^{st}(\mathbf{u})\boldsymbol{\varepsilon}(\mathbf{u})^\top] = \mathrm{E}[\mathbf{x}^{st}(\mathbf{u})]\mathrm{E}[\boldsymbol{\varepsilon}(\mathbf{u})^\top] = \mathbf{0} = \mathrm{E}[\boldsymbol{\varepsilon}(\mathbf{u})]\mathrm{E}[\mathbf{x}^{st}(\mathbf{u})^\top] = \mathrm{E}[\boldsymbol{\varepsilon}(\mathbf{u})\mathbf{x}^{st}(\mathbf{u})^\top].$$

Finally, $\mathrm{E}[\boldsymbol{\varepsilon}(\mathbf{u})\boldsymbol{\varepsilon}(\mathbf{u})^\top] = \mathrm{Cov}(\boldsymbol{\varepsilon}) = \mathbf{I}_r$, and thus,

$$\mathbf{I}_{p+r} - \mathrm{Cov}(\mathbf{x}^{st*}) = \mathbf{I}_{p+r} - \begin{pmatrix} \mathrm{Cov}(\mathbf{x}^{st}) & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_r \end{pmatrix} = \begin{pmatrix} \mathbf{I}_p - \mathrm{Cov}(\mathbf{x}^{st}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

Since we are only summing scaled squares ($A^2 = AA^\top$) of matrices of this form, the resulting matrix, $\mathbf{M}_{\mathrm{var}}$ is of the same form. The arguments for the other scatter matrices follows similarly.

Let us now take a closer look at the eigenvectors of $\mathbf{M}(\mathbf{x}^{st*})$, denoted by

$$\mathbf{v}_i^* = \begin{pmatrix} \mathbf{v}_i \\ \mathbf{v}_i^{\mathrm{AUG}} \end{pmatrix},$$

where $\mathbf{v}_i \in \mathbb{R}^p$, $\mathbf{v}_i^{\mathrm{AUG}} \in \mathbb{R}^r$, and $\mathbf{v}_i^*$ corresponds to the eigenvalue d_i^* of $\mathbf{M}(\mathbf{x}^{st*})$. Arrange the eigenvalues such that

$$d_1^* \geq d_2^* \geq \cdots \geq d_{p+r}^*.$$

The eigenpair satisfies

$$\mathbf{M}(\mathbf{x}^{st*})\mathbf{v}_i^* = \begin{pmatrix} \mathbf{M}(\mathbf{x}^{st}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{v}_i \\ \mathbf{v}_i^{\mathrm{AUG}} \end{pmatrix} = \begin{pmatrix} \mathbf{M}(\mathbf{x}^{st})\mathbf{v}_i \\ \mathbf{0} \end{pmatrix} = d_i^*\mathbf{v}_i^* = \begin{pmatrix} d_i^*\mathbf{v}_i \\ d_i^*\mathbf{v}_i^{\mathrm{AUG}} \end{pmatrix}.$$

This shows that $d_i^*\mathbf{v}_i^{\mathrm{AUG}} = \mathbf{0}$, for all i, and gives us

$$|d_i^*|\|\mathbf{v}_i^{\mathrm{AUG}}\| = 0.$$

For all $i \leq q$, the eigenvalues d_i^* are non-zero, and hence it follows that the augmented part of the eigenvector has the norm $\|\mathbf{v}_i^{\mathrm{AUG}}\| = 0$. However, this may not hold for $i > q$, since $d_i^* = 0$. Hence, the norms of the augmented components $\mathbf{v}_i^{\mathrm{AUG}}$ provide insight into the rank q of $\mathbf{M}(\mathbf{x}^{st})$. Using the information provided by the eigenvectors and eigenvalues, we present three methods for rank estimation.

Rank estimation via augmented eigenvectors. Let $\mathbf{v}_0^{\mathrm{AUG},j} := \mathbf{0}$ and define the function

$$f : \{0, 1, \ldots, p\} \to \mathbb{R}, \quad f(i) = \frac{1}{s}\sum_{j=0}^{s} \|\mathbf{v}_i^{\mathrm{AUG},j}\|^2,$$

where s is the number of repetitions of the augmentation procedure to average out the randomness for simulating the white noise fields.

From the properties of the augmented parts of the eigenvectors, we know that $f(i)$ should be approximately zero for $i \leq q$. Hence the point at which f starts increasing gives us an estimate of the dimension.

Normalized scree plot. Define $d_{p+1} := 0$, and following Luo and Li (2021), define the normalized scree function by

$$\Phi : \{0, 1, \ldots, p\} \to \mathbb{R}, \quad \Phi(l) = \frac{d_{l+1}}{\sum_{i=1}^{l+1} d_i}.$$

The scree function starts at value $\Phi(0) = 1$, and goes down to zero once $i > q$. So the sufficiently large non-zero value of the scree function gives an estimate of q. Note that this estimate assumes that there is at least one nonstationary signal. If all signals were to be stationary, all of the eigenvalues d_i would be zero and the values of $\Phi(\ell)$ would not be defined for any ℓ. The problem arises from the division by zero. This can be easily mitigated by replacing the denominator by $1 + \sum_{i=1}^{l+1} d_i$. However, it should be kept in mind that doing so breaks the scale invariance of the estimate. For our purposes, we can ignore this issue as our model assumes by default that there is at least one nonstationary signal.

Augmentation estimator. The final estimator combines both the eigenvalue and eigenvector information:

$$g : \{0, 1, \ldots, p\} \to \mathbb{R}, \quad g(k) = \Phi(k) + \sum_{i=1}^{k} f(i),$$

and the estimator of the number of signal components is given by

$$\hat{q} = \underset{k \in \{0,\ldots,p\}}{\arg\min}\ g(k).$$

In the ideal case, the scree function is non-zero and decreasing up to q. For those values, f is zero. This relation flips after q, at which point f will be non-zero and increasing, and the scree function vanishes. Consequently, the two functions intersect at q producing a cusp in g. This point defines the augmentation estimator.

In practice, rank estimation based on augmented eigenvalues or normalized scree plots typically requires graphical inspection, with the change point selected visually. This subjective step may lead to ambiguity and a lack of reproducibility. In contrast, plotting the augmentation estimator produces a so-called ladle plot, which ideally exhibits a clear minimum and therefore yields a unique and objective estimate of the rank. For this reason, we focus on the augmentation estimator in the simulation studies presented below.

The three estimators introduced above are formulated for settings that can be expressed within a generalized eigenvalue–eigenvector framework. When estimators are combined, as in the case of `spSSAcomp`, the problem is instead cast as one of joint diagonalization. Nevertheless, the augmentation procedure remains applicable: in this setting, the eigenvalues are replaced by sums of the absolute values of the corresponding pseudo-eigenvalues.

5 Simulations II

In the following simulations, we study the finite sample performance of the augmentation estimator. We use Setting 4 from Section 3 with the side length $\ell = 60$ and the sample size $n = 3600$. We use a 4-by-4 grid partitioning of the domain. We test the method for varying numbers of dimensions for the augmented data, namely, for $r = 1, 5, 10, 15$. For all dimensions of the augmented data, we perform the augmentation estimation procedure for 2000 repetitions. Note that the true nonstationary dimension is $q = 3$. We then compute how many times each result appeared and measure the appearance rate of each computed value. The results are presented as proportions of the horizontal bar with different values being assigned different colors. The correct value is represented by the color gray. Underestimations are represented by blue and overestimates by red. The darker the color, the more we have over, or, underestimated. We fix the number of trials s for the augmented vectors to $s = 10$. This is a standard assumption in the literature and is done for example in Luo and Li (2021); Radojičić and Nordhausen (2024).

The results of the simulations can be seen in Figure 8. For $r = 1$, most methods tend to overestimate. As we increase the number of augmented signals, each method starts to underestimate. This happens because the "mass" of the signals falls on the augmented part which is purely stationary.

The results reflect what was expected: SPSSASIR finds one signal, its own, SPSSASAVE and SPSSACOR find two signals, and the jointly diagonalized estimate correctly estimates the nonstationary dimension. We see that SPSSACOMB gets the best results for $r = 5$ and $r = 10$. Thus, we recommend r be in the range $[5, 10]$ as a suitable choice of parameters. This is in line with other recommendations in the literature, see for example Radojičić and Virta (2025), who show that for too large r, a high-dimensional regime might be obtained which, if not properly handled, induces bias when estimating the eigenvalues. Section B shows similar plots for the different settings given in Section 3.

6 Spatial stationary subspace analysis for the Kola moss data

In this section, we apply spSSA to the moss data from the Kola project and estimate the dimension of the underlying stationary subspace. The data are publicly available in the `R` package `StatDA` Filzmoser (2023) and are described in detail in Reimann et al. (1998). The same data set was also analyzed as a real-data example in Nordhausen et al. (2015) under the assumption that it is stationary.

From the full data set out of 39 elements, we consider all elements without missing values. Thus, for our analysis, we include the elements Ag, Al, As, B, Ba, Be, Bi, Ca, Cd, Co, Cr, Cu, Fe, Hg, K, La, Mg, Mn, Mo, Na, Ni, P, Pb, Rb, S, Sb, Sc, Se, Si, Sr, Th, Tl, U, V, Y, and Zn, while excluding the elements Au, Pd, and Pt. Consequently, the resulting random field $\mathbf{x}(\mathbf{u})$ is 36-dimensional. The random field was measured in total at 594 locations.

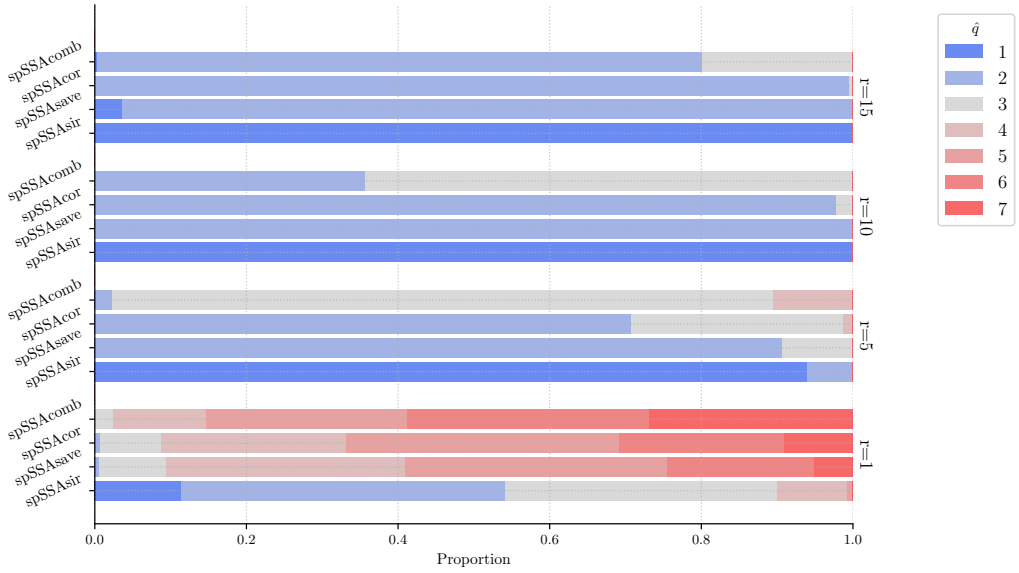

Figure 8: Rank estimation for Setting 4 using augmented noise dimensions $r = 1, 5, 10, 15$. The simulation is done so that each signal has 3600 data points and the mixed signals are processed using a $(4, 4)$ partition of the domain. Neutral gray indicates successful estimation. Blue is used to indicate underestimation and red overestimation. The darker the color, the worse the estimation.

The observed data are compositional in nature and therefore require an appropriate transformation before spSSA can be applied. Following the approach of Nordhausen et al. (2015), to which we also refer for further details, we consider $\mathbf{x}_{\text{ilr}}$, the *isometric log-ratio* (ilr) transformation of the data. Since the ilr transformation maps a D-part composition to $\mathbb{R}^{D-1}$, this transformation reduces the dimension by one, yielding a 35-dimensional random field.

We apply our preferred method, SPSSACOMB, which requires the specification of a kernel function and a partition of the spatial domain. We choose a ball kernel with radius 50 km, matching the choice in Nordhausen et al. (2015). The partition of the domain is shown in Figure 9. Note that we are again using the scaled version of SPSSACOR as a subroutine of SPSSACOMB.

In Figure 9, the spatial domain is displayed twice. The left panel shows the measurement locations together with an illustration of the ball kernel with radius 50 km. The right panel displays the partition used in the analysis, demonstrating that the proposed methodology is applicable also for irregularly shaped partitions.

We now estimate the dimension $\hat{q}$ of the nonstationary subspace of $\mathbf{x}_{\text{ilr}}$. To this end, we apply SPSSACOMB using the kernel function and the partition described above and set the tuning parameters to $r = 10$ and $s = 10$. The three criteria for rank estimation, $f(k)$, $\Phi(k)$, and $g(k)$, are visualized in Figure 10. The minimum of $g(k)$ in the ladle plot is highlighted in blue, providing the estimate $\hat{q} = 5$.

While the estimators $f(k)$ and $\Phi(k)$ individually do not indicate a clear choice for $\hat{q}$, the combined criterion $g(k)$ exhibits a well-defined minimum, illustrating for example the advantage of the ladle plot over a scree plot.

Based on the estimated rank $\hat{q} = 5$, the unmixing matrices $\mathbf{W}_{\text{comb},n}$ and $\mathbf{W}_{\text{comb},s}$ were constructed. Using these matrices, we separate the two subspaces and visualize all five nonstationary signals as well as one randomly chosen representative stationary signal $\mathbf{s}_{21}$. The resulting spatial patterns are shown in Figure 11. In the figure, the values of the extracted signals are displayed using symbols corresponding to percentile ranges. We also give the pseudo-eigenvalues of the nonstationary signals and the representative stationary signal for all $\mathbf{M}$ matrices involved in the joint diagonalization in Table 2.

The results indicate a clear contrast between the two components: the stationary signal exhibits a fairly uniform spatial distribution across all value ranges, whereas the nonstationary signals display pronounced regional structures and spatial trends. The pseudo-eigenvalues of the individual matrices employed in SPSSACOMB provide additional insight into

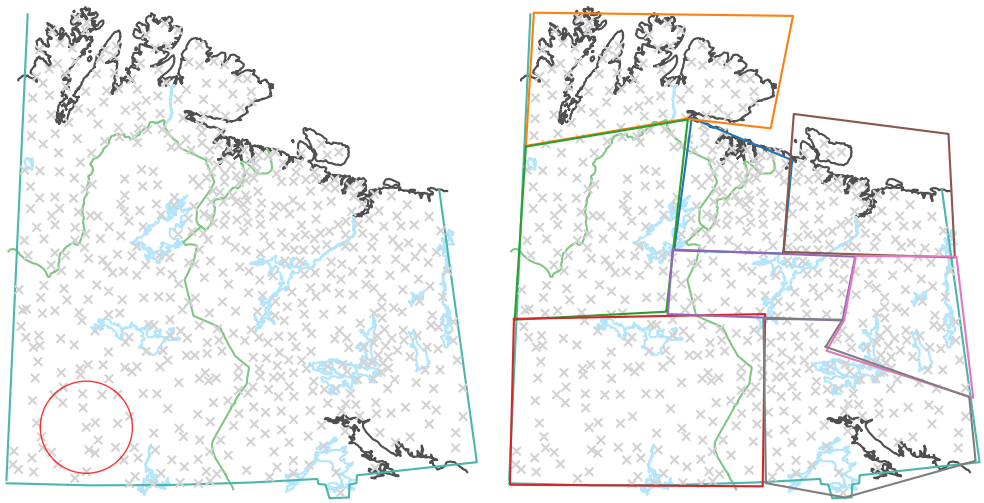

Figure 9: Two maps of the measurement locations for the Kola data. The measurement locations are marked with a gray $\times$ symbol. On the left, the red circle shows a ball of radius 50km. On the right, the used partition is shown.

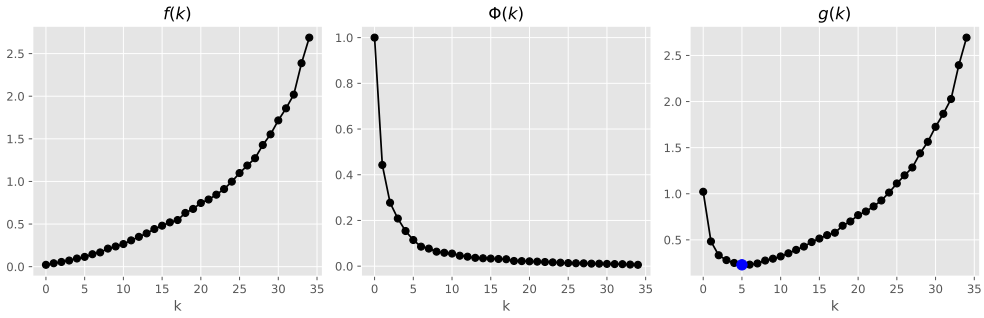

Figure 10: Plot of rank estimators f, Φ, g on the Kola data with the estimated rank $\hat{q} = 5$, highlighted in blue on the g ladle.

Table 2: The pseudo-eigenvalues of the $\mathbf{M}$ matrices from approximate joint diagonalization for the $\hat{q} = 5$ nonstationary signals and the representative stationary signal $\mathbf{s}_{21}$. Note that since there are 5 nonstationary signals, the 21st stationary signal corresponds to the 26th pseudo-eigenvalue.

	d_1	d_2	d_3	d_4	d_5	d_{26}
$\mathbf{M}_{\text{mean}}$	0.025	0.139	0.766	0.622	0.054	0.107
$\mathbf{M}_{\text{var}}$	3.251	1.899	0.869	0.960	1.639	0.486
$\mathbf{M}^{s}_{\text{cor}}$	3.111	3.038	2.772	2.596	1.950	0.471

the nature of the nonstationarity present in the extracted signals. We report these values for the five nonstationary components in Table 2. An inspection of the pseudo-eigenvalues suggests that the dominant sources of nonstationarity arise from spatial dependence and spatially varying variance, with only a minor contribution being attributable to spatial location trends.

7 Conclusion

In this paper, we developed tools for spatial stationary subspace analysis. We showed via simulations that given the dimension of the nonstationary subspace, we can successfully separate the two subspaces. If only a single type of

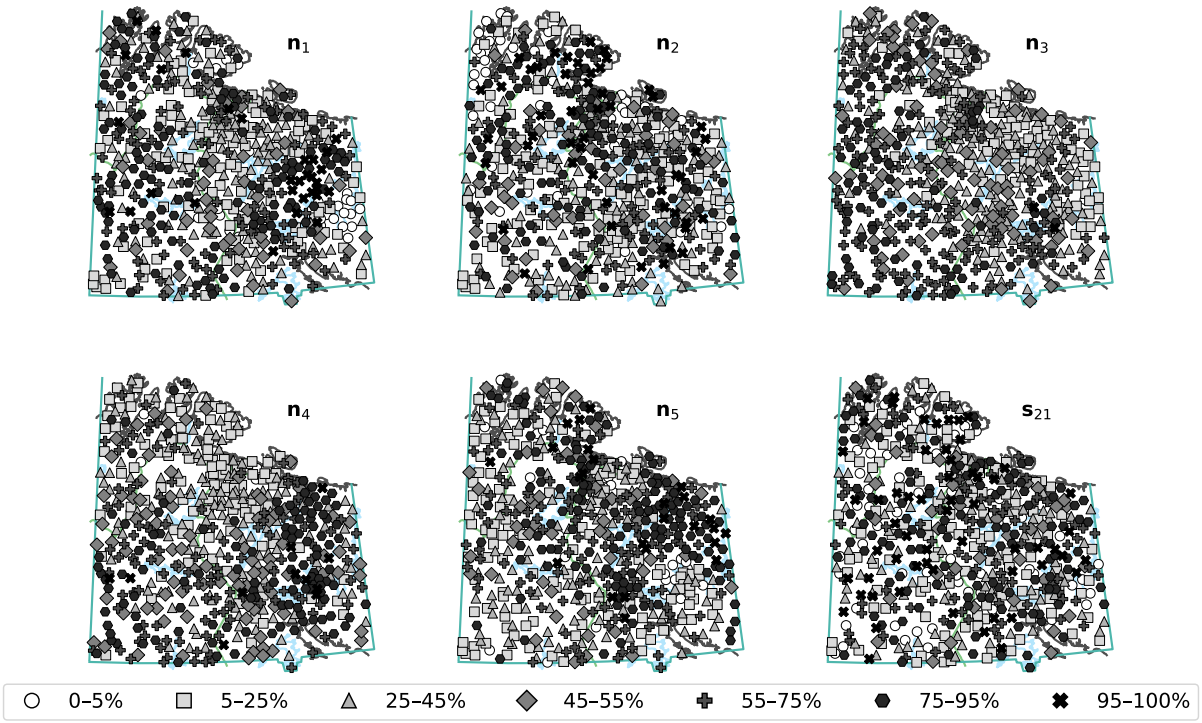

Figure 11: Nonstationary components and one stationary component of $\mathbf{x}_{\text{ilr}}$. Symbols indicate percentile ranges of values.

nonstationarity is present, the method developed specifically for such a case performs best with SPSSACOMB being a close second. If multiple types of nonstationarities are present, SPSSACOMB clearly achieves the best performance. It is hence our recommendation to use SPSSACOMB. The importance of a proper partition was also demonstrated. It is recommended to use a granular partition so that all parts still have enough data points to do proper estimations.

In any SSA method, the estimation of the subspace dimensions is of central practical importance. In this paper, we propose three estimators: one based on information contained in the eigenvectors, one relying on the eigenvalues, and a third estimator that combines both sources of information.

The combined estimator is closely related to the ladle estimator originally introduced for i.i.d. data and stationary time series in Luo and Li (2016); Nordhausen and Virta (2018). A key difference is that, in its original formulation, the eigenvector-based component relies on bootstrapping. While feasible in i.i.d. and time series settings, such resampling becomes substantially more computationally demanding in a spatial framework. The approach proposed here avoids this difficulty and is therefore better suited to spatial data analysis.

It is also worth noting that the stationary subspace analysis (SSA) methods for time series discussed in Flumian et al. (2024) share a structural similarity with the spSSA methods considered in this paper. Consequently, the dimension estimation strategy developed here can be directly transferred to these SSA methodologies with minimal adaptation.

The second simulation study demonstrates that the augmentation estimator—combining information from both eigenvalues and eigenvectors—successfully recovers the dimension of the nonstationary subspace. In addition, the results consistently indicate that SPSSACOMB provides the most reliable performance across scenarios. Our recommendations for tuning parameter choices are consistent with existing literature: specifically, we suggest fixing $s = 10$ and selecting r from the range $[5, 10]$.

All in all, the proposed methodology equips practitioners with a practical and robust set of tools for separating multivariate spatial data into stationary and nonstationary components. These components can subsequently be modeled independently using appropriate techniques, which is typically far more tractable than analyzing the original high-

dimensional data in which both structures are confounded. This practical value is illustrated through the analysis of the Kola moss data set, where a 35-dimensional random field is effectively decomposed into a five-dimensional nonstationary component and a 30-dimensional stationary component.

Acknowledgments

We acknowledge computational resources from CSC – IT Center for Science, Finland. PS and KN were supported by the Research Council of Finland (363261). AMR acknowledges funding from the French National Research Agency (ANR) under the Investissements d'Avenir program (ANR-17-EURE-0010).

A Comparing SPSSACOR to scaled SPSSACOR

To demonstrate the comparability in terms of performance for the SPSSACOR methods, we repeat the simulations of Setting 3 from Section 3. This time, we only include two curves in the resulting plots: one for the unscaled SPSSACOR, denoted default, and on for the scaled version, denoted scaled. The results can be seen in Figure 12. The results demonstrate clearly that the scaled SPSSACOR causes no loss in performance, and in fact performs slightly better than its unscaled counterpart.

Let us also demonstrate the need to use the scaled version by comparing norms of the $\mathbf{M}$ matrices. Consider the scenario where the 5-variate stationary signal $\mathbf{s}(\mathbf{u})$ from Section 3 on 2500 data points is mixed with some orthogonal matrix $\mathbf{A} \in \mathbb{R}^{5\times 5}$ to get $\mathbf{x}(\mathbf{u}) = \mathbf{A}\mathbf{s}(\mathbf{u})$. Consider further the standardized data $\mathbf{x}^{st}(\mathbf{u})$. Let us now compute $\mathbf{M}_{\text{mean}}$, $\mathbf{M}_{\text{var}}$, $\mathbf{M}_{\text{cor},f}$ and $\mathbf{M}^s_{\text{cor},f}$ from $\mathbf{x}^{st}(\mathbf{u})$. For the kernel function f, we use the ball kernel with radius 2.2. In Table 3, we have tabulated the average norm of these $\mathbf{M}$ matrices over 1000 repetitions using three different k-by-k partitions for $k = 2, 3, 4$. This data shows that the norms of all $\mathbf{M}$ matrices increase as the number of partitions increases, but the $\mathbf{M}_{\text{cor},f}$ matrices differ by two orders of magnitude from the other $\mathbf{M}$ matrices. Even though $\mathbf{M}^s_{\text{cor},f}$ is still of greater magnitude than $\mathbf{M}_{\text{mean}}$ and $\mathbf{M}_{\text{var}}$, it is at least of comparable scale. Note that we chose to demonstrate this effect on stationary signals since the mean and variance estimators are close to zero, making the need for scaling more apparent.

Table 3: Average norms of $\mathbf{M}$ matrices for a mixed stationary 5-variate standardized random field.

Partition	$\mathbf{M}_{\text{mean}}$	$\mathbf{M}_{\text{var}}$	$\mathbf{M}^s_{\text{cor},f}$	$\mathbf{M}_{\text{cor},f}$
$(2, 2)$	0.028	0.044	0.084	5.364
$(3, 3)$	0.056	0.105	0.193	11.727
$(4, 4)$	0.092	0.181	0.323	18.274

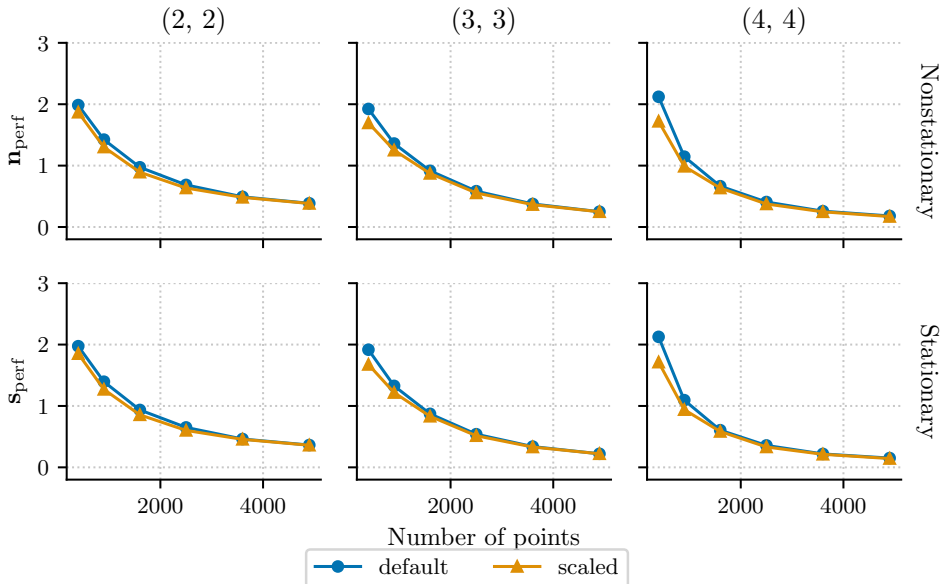

Figure 12: Comparison between SPSSACOR (default) and scaled SPSSACOR (scaled) on data from Setting 3 from Section 3.

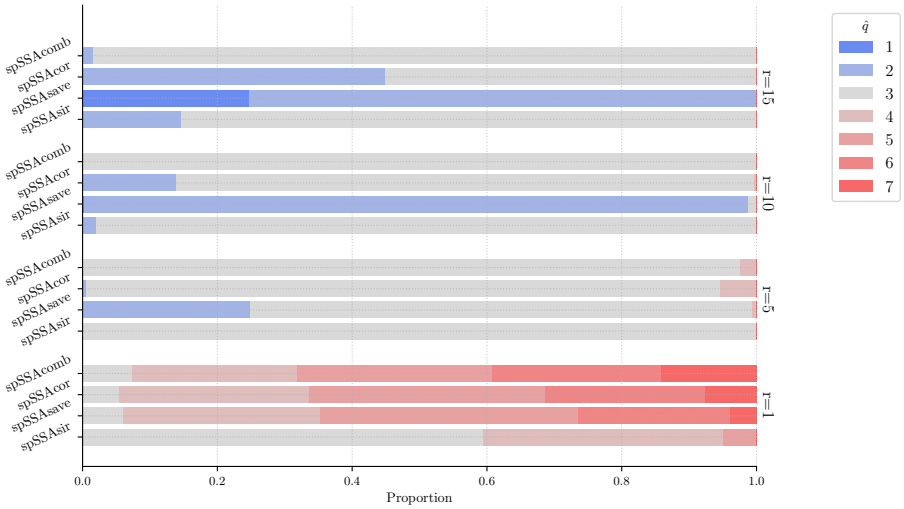

Figure 13: Rank estimation for Setting 1 using augmented noise dimensions $r = 1, 5, 10, 15$. The simulation is done so that each signal has 3600 data points and the mixed signals are processed using a $(4, 4)$ partition of the domain. Neutral gray indicates successful estimation. Blue is used to indicate underestimation and red overestimation. The darker the color, the worse the estimation.

B Additional simulations for the augmentation estimator

In this section, we repeat the simulation from Section 5 for Settings 1, 2, and 3 from Section 3. Everything is kept the same except the types of nonstationary signals. The results for these three settings can be seen in Figure 13, Figure 14, and Figure 15, respectively. The results show that if only one type of nonstationarity is present, we can determine the order even more accurately. We further see that the recommended values of r work in all settings.

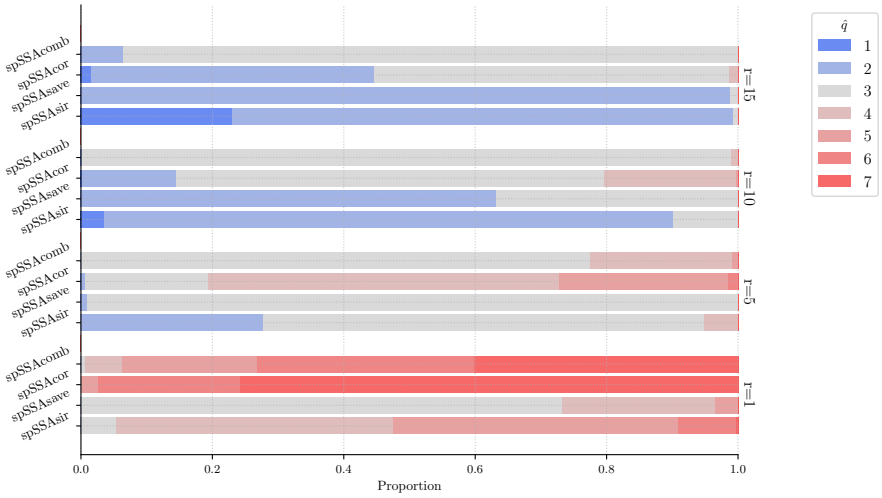

Figure 14: Rank estimation for Setting 2 using augmented noise dimensions $r = 1, 5, 10, 15$. The simulation is done so that each signal has 3600 data points and the mixed signals are processed using a $(4, 4)$ partition of the domain. Neutral gray indicates successful estimation. Blue is used to indicate underestimation and red overestimation. The darker the color, the worse the estimation.

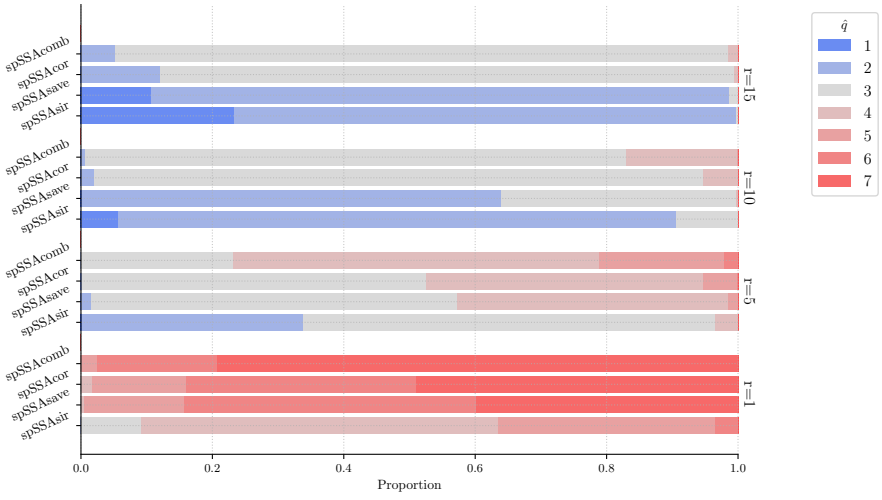

Figure 15: Rank estimation for Setting 3 using augmented noise dimensions $r = 1, 5, 10, 15$. The simulation is done so that each signal has 3600 data points and the mixed signals are processed using a $(4, 4)$ partition of the domain. Neutral gray indicates successful estimation. Blue is used to indicate underestimation and red overestimation. The darker the color, the worse the estimation.

References

Anderes, E. B. and Stein, M. L. (2011). Local likelihood estimation for nonstationary random fields. *Journal of Multivariate Analysis*, 102(3):506–520.

Apanasovich, T. V., Genton, M. G., and Sun, Y. (2012). A valid Matérn class of cross-covariance functions for multivariate random fields with any number of components. *Journal of the American Statistical Association*, 107(497):180–193.

Bachoc, F., Genton, M. G., Nordhausen, K., Ruiz-Gazen, A., and Virta, J. (2020). Spatial blind source separation. *Biometrika*, 107(3):627–646.

Clarkson, D. (1988). A least squares version of algorithm AS 211: The F-G diagonalization algorithm. *Journal of the Royal Statistical Society, Series C*, 37:317–321.

Cook, R. (2000). SAVE: A method for dimension reduction and graphics in regression. *Communications in Statistics - Theory and Methods*, 29:2109–2121.

Filzmoser, P. (2023). *StatDA: Statistical Analysis for Environmental Data*. R package version 1.7.11.

Flumian, L., Matilainen, M., Nordhausen, K., and Taskinen, S. (2024). Stationary subspace analysis based on second-order statistics. *Journal of Computational and Applied Mathematics*, 436:115379.

Fotheringham, A. S. and Wong, D. W. S. (1991). The modifiable areal unit problem in multivariate statistical analysis. *Environment and Planning A: Economy and Space*, 23(7):1025–1044.

Fuentes, M. (2002). Spectral methods for nonstationary spatial processes. *Biometrika*, 89(1):197–210.

Gelfand, A. E., Schmidt, A. M., Banerjee, S., and Sirmans, C. F. (2004). Nonstationary multivariate process modeling through spatially varying coregionalization. *Test*, 13:263–312.

Genton, M. G. and Kleiber, W. (2015). Cross-covariance functions for multivariate geostatistics. *Statistical Science*, 30(2):147–163.

Gneiting, T., Kleiber, W., and Schlather, M. (2010). Matern cross-covariance functions for multivariate random fields. *Journal of the American Statistical Association*, 105:1167–1177.

Goulard, M. and Voltz, M. (1992). Linear coregionalization model: Tools for estimation and choice of cross-variogram matrix. *Mathematical Geology*, 24:269–286.

Guttorp, P. and Gneiting, T. (2006). Studies in the History of Probability and Statistics XLIX on the Matérn Correlation Family. *Biometrika*, 93(4):989–995.

Hara, S., Kawahara, Y., Washio, T., and von Bünau, P. (2010). Stationary subspace analysis as a generalized eigenvalue problem. In Wong, K. W., Mendis, B. S. U., and Bouzerdoum, A., editors, *Neural Information Processing. Theory and Algorithms*, pages 422–429, Berlin, Heidelberg. Springer Berlin Heidelberg.

Ingebrigtsen, R., Lindgren, F., and Steinsland, I. (2014). Spatial models with explanatory variables in the dependence structure. *Spatial Statistics*, 8:20–38.

Kleiber, W. and Nychka, D. (2012). Nonstationary modeling for multivariate spatial processes. *Journal of Multivariate Analysis*, 112:76–91.

Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86:316–327.

Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society, Series B*, 73(4):423–498.

Liski, E., Nordhausen, K., Oja, H., and Ruiz-Gazen, A. (2016). Combining linear dimension reduction subspaces. In Agostinelli, C., Basu, A., Filzmoser, P., and Mukherjee, D., editors, *Recent Advances in Robust Statistics: Theory and Applications*, pages 131–149, New Delhi. Springer India.

Luo, W. and Li, B. (2016). Combining eigenvalues and variation of eigenvectors for order determination. *Biometrika*, 103:875–887.

Luo, W. and Li, B. (2021). On order determination by predictor augmentation. *Biometrika*, 108:557–574.

Miettinen, J., Nordhausen, K., and Taskinen, S. (2017). Blind source separation based on joint diagonalization in R: The packages JADE and BSSasymp. *Journal of Statistical Software*, 76:1–31.

Muehlmann, C., Bachoc, F., and Nordhausen, K. (2022). Blind source separation for non-stationary random fields. *Spatial statistics*, 47.

Muehlmann, C., Bachoc, F., Nordhausen, K., and Yi, M. (2024). Test of the latent dimension of a spatial blind source separation model. *Statistica Sinica*, 34:837–865.

Nordhausen, K., Oja, H., Filzmoser, P., and Reimann, C. (2015). Blind source separation for spatial compositional data. *Mathematical Geosciences*, 47(7):753–770.

Nordhausen, K. and Ruiz-Gazen, A. (2022). On the usage of joint diagonalization in multivariate statistics. *Journal of Multivariate Analysis*, 188:104844.

Nordhausen, K. and Virta, J. (2018). Ladle estimator for time series signal dimension. In *2018 IEEE Statistical Signal Processing Workshop (SSP)*, pages 428–432.

Openshaw, S. (1984). The modifiable areal unit problem. *Concepts and techniques in modern geography*.

Paciorek, C. J. and Schervish, M. J. (2006). Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics*, 17(5):483–506.

Piccolotto, N., Bögl, M., Muehlmann, C., Nordhausen, K., Filzmoser, P., and Miksch, S. (2022). Visual parameter selection for spatial blind source separation. *Computer Graphics Forum*, 41(3):157–168.

Radojičić, U. and Nordhausen, K. (2024). Order determination in second-order source separation models using data augmentation. In Ansari, J., Fuchs, S., Trutschnig, W., Lubiano, M. A., Gil, M. Á., Grzegorzewski, P., and Hryniewicz, O., editors, *Combining, Modelling and Analyzing Imprecision, Randomness and Dependence*, pages 371–379, Cham. Springer Nature Switzerland.

Radojičić, U., Lietzén, N., Nordhausen, K., and Virta, J. (2021). Dimension estimation in two-dimensional pca. In *2021 12th International Symposium on Image and Signal Processing and Analysis (ISPA)*, pages 16–22.

Radojičić, U., Lietzén, N., Nordhausen, K., and Virta, J. (2025). Order determination for tensor-valued observations using data augmentation. *Journal of Computational and Graphical Statistics*, pages 1–11.

Radojičić, U. and Virta, J. (2025). Dimension estimation in a spiked covariance model using high-dimensional data augmentation. *Biometrika*, 112(4):asaf052.

Reimann, C., Äyräs, M., Chekushin, V. A., Bogatyrev, I. V., Boyd, R., Caritat, P. d., Dutter, R., Finne, T. E., Halleraker, J. H., Jæger, Ø., et al. (1998). *Environmental geochemical atlas of the central Barents region.* NGU-GTK-CKE Special Publication, Geological Survey of Norway, Trondheim.

Sampson, P. D. (2010). Constructions for nonstationary spatial processes. In Gelfand, A. E., Diggle, P. J., Fuentes, M., and Guttorp, P., editors, *Handbook of Spatial Statistics*, pages 119–130. CRC Press, Boca Raton.

Sampson, P. D. and Guttorp, P. (1992). Nonparametric estimation of nonstationary spatial covariance structure. *Journal of the American Statistical Association*, 87(417):108–119.

Sipilä, M., Nordhausen, K., and Taskinen, S. (2024). Nonlinear blind source separation exploiting spatial nonstationarity. *Information Sciences*, 665:120365.

Tobler, W. R. (1970). A computer movie simulating urban growth in the detroit region. *Economic Geography*, 46(Suppl.):234–240.

von Bünau, P., Meinecke, F. C., Király, F. C., and Müller, K.-R. (2009). Finding stationary subspaces in multivariate time series. *Physical Review Letters*, 103:214101.

Vu, Q., Zammit-Mangion, A., and Cressie, N. (2022). Modeling nonstationary and asymmetric multivariate spatial covariances via deformations. *Statistica Sinica*, 32(4):2071–2093.

Wackernagel, H. (2003). *Multivariate Geostatistics*. Springer.