



Robust signal dimension estimation via SURE

Joni Virta¹ · Niko Lietzén¹ · Henri Nyberg¹

Received: 31 October 2022 / Revised: 11 July 2023
© The Author(s) 2023

Abstract

The estimation of signal dimension under heavy-tailed latent variable models is studied. As a primary contribution, robust extensions of an earlier estimator based on Gaussian Stein's unbiased risk estimation are proposed. These novel extensions are based on the framework of elliptical distributions and robust scatter matrices. Extensive simulation studies are conducted in order to compare the novel methods with several well-known competitors in both estimation accuracy and computational speed. The novel methods are applied to a financial asset return data set.

Keywords Elliptical distribution · Principal component analysis · Risk estimate · Scatter matrix

1 Introduction

1.1 Premise

Modern data sets exhibit increasingly large sizes and often the first step in analysing them is the implementation of suitable dimension reduction procedures, for example, principal component analysis (PCA). A fundamental question pertaining to any such reduction is the choosing of the correct amount of latent components to retain: underestimating their amount leads to losing important information, whereas picking too many components inevitably leads in the later stages of the analysis to modelling of noise and increased computational burden that could have been avoided with a more careful choice of the dimension.

✉ Joni Virta
joni.virta@utu.fi
Niko Lietzén
niko.lietzen@utu.fi
Henri Nyberg
henri.nyberg@utu.fi

¹ Department of Mathematics and Statistics, University of Turku, Turku, Finland

Besides being large in size and volume, in many applications, such as economics and finance, it is common for the data sets to display heavier tails than possessed by the Gaussian distribution. Consider, for example, stock market returns which are often modelled with distributions having infinite variance (Borak et al. 2011). In such cases, the estimation of the latent dimension in dimension reduction is further complicated as one cannot rely any more on the covariance matrix, on whose eigenvalues many of the standard dimension estimation methods rely (a review is given later in this section).

With the above scenario in our mind, the objective of the current paper is to study and develop dimension estimation in the context of multivariate data sets exhibiting arbitrarily heavy tails. We will base our theoretical framework on the concepts of elliptical family and Stein’s unbiased risk estimation, elaborated in more detail next.

1.2 Elliptical latent variable model

We assume that our data x_1, \dots, x_n is an i.i.d. sample of p -variate vectors generated from the elliptical latent variable model

$$x_i = \mu + VDz_i, \tag{1}$$

where $\mu \in \mathbb{R}^p$, $V \in \mathbb{R}^{p \times p}$ is an orthogonal matrix, z_i obeys a spherical distribution (Fang 2018), i.e., $z_i \sim Oz_i$ for any orthogonal matrix $O \in \mathbb{R}^{p \times p}$, and $D \in \mathbb{R}^{p \times p}$ is a diagonal matrix with the diagonal elements $\sigma_1 \geq \dots \geq \sigma_d > \sigma = \dots = \sigma$ with $\sigma > 0$. Conceptually, the model says that the observed x_i are obtained by mixing the principal components Dz_i with the matrix V and by applying a location shift μ . The final $p - d$ principal components in Dz_i are orthogonally invariant, meaning that they are essentially “structureless” and, as is typical, we view them as *noise*. Thus the main objective in the model (1) is to estimate the latent *signals*, i.e., the first d elements of Dz_i along with their number d . Note that while the scales of D and z_i are confounded in the model (1), this does not matter as long as one considers only the compound quantity Dz_i .

The model (1) gets still more intuitive form in the special case where z_i obeys the standard Gaussian distribution, the most well-known example of an elliptical distribution. In this case model (1) reduces to the probabilistic PCA model (Tipping and Bishop 1999)

$$x_i = \mu + V_0y_i + \varepsilon_i, \tag{2}$$

where the loading matrix $V_0 \in \mathbb{R}^{p \times d}$ contains the d first columns of V as its columns,

$$y_i \sim \mathcal{N}_d\{0, \text{diag}(\sigma_1^2 - \sigma^2, \dots, \sigma_d^2 - \sigma^2)\},$$

and $\varepsilon_i \sim \mathcal{N}_p(0, \sigma^2 I_p)$ is independent of y_i . Model (2) reveals that, in the Gaussian case, the d -dimensional signal residing in the column space of V_0 is explicitly corrupted with the noise vectors ε_i , and the model can be seen as a factor model (additive contamination of latent signals with noise). However, this intuitive representation does

not apply to any other elliptical distribution (that is, (1) cannot in general be written in the form (2) for elliptical z_i) and, hence, (1) is usually viewed in the literature as an elliptical PCA model or as an elliptical subsphericity model, rather than as a factor model. Naturally, it would also be possible to consider an elliptical variant of the factor model (2) where both y_i and ε_i are assumed to be non-Gaussian, see, for example, Pison et al. (2003).

The standard method of extracting the factors z_i (or the corresponding subspace) in the Gaussian model (2) is through PCA. Namely, one computes the first d eigenvectors of the covariance matrix of x_i and projects the observations onto their span. However, the success of this procedure hinges crucially on the knowledge of the latent dimension $d < p$, usually unknown to us in practice. As the misspecification of the dimension has ill consequences in practice (either missing part of the signal or riddling our estimates of the latent factors with noise), an important part of solving the factor model (2) is the accurate estimation of d . We next review a particular method for accomplishing this under the model (2), on which our subsequent developments are also based.

1.3 Stein's unbiased risk estimate

In Ulfarsson and Solo (2015), the latent dimension d in Gaussian PCA was estimated through the application of the *Stein's unbiased risk estimate* (SURE), a general technique for determining the optimal values of tuning parameters (of which the latent dimension d is an example) of estimation procedures.

Ignoring the model (2) for a moment, we briefly recall the basic idea behind the SURE: Assume that we observe the i.i.d. univariate w_1, \dots, w_n , generated as $w_i = a_i + e_i$, where $a_i \in \mathbb{R}$ are constant and the errors satisfy $e_i \sim \mathcal{N}(0, \tau^2)$ for some $\tau^2 > 0$. Assume further that for each a_i we have a corresponding estimator $\hat{a}_i(w)$, viewed as a (differentiable) function of the data $w = (w_1, \dots, w_n)$. Then, the SURE R corresponding to the estimators a_i is defined as

$$R = \frac{1}{n} \sum_{i=1}^n \{w_i - \hat{a}_i(w)\}^2 + \frac{2\tau^2}{n} \sum_{i=1}^n \frac{\partial}{\partial w_i} \hat{a}_i(w) - \tau^2. \quad (3)$$

In their celebrated paper (Stein 1981), Stein proved that $E(R) = (1/n) \sum_{i=1}^n E\{\hat{a}_i(w) - a_i\}^2$, showing that R is an unbiased estimator of the risk associated with the estimators \hat{a}_i , in the complete absence of the true means a_i . A multivariate version of SURE was in Ulfarsson and Solo (2015) adapted to the PCA model (2) (conditional on the y_i) to estimate the expected risk associated with any particular choice of the latent dimension d . The estimate of d is then chosen to be the dimensionality for which the risk is minimized. See also the earlier work (Ulfarsson and Solo 2008).

The obtained estimator was in Ulfarsson and Solo (2015) shown to be highly successful under the Gaussian model (2). However, it is clear that the estimator cannot obtain the same level of efficiency under the wider elliptical model (1). There are two reasons for this: (i) the SURE-criterion was in Ulfarsson and Solo (2015) derived

strictly under the Gaussian assumption and, more importantly, (ii) many standard elliptical distributions (e.g., multivariate Cauchy distribution) do not have enough finite moments so that the covariance matrix, on which the SURE-criterion is based, would even exist. Hence, the estimator of Ulfarsson and Solo (2015) is not even necessarily well-defined under the elliptical model (1) (on the population level)!

1.4 The scope of the current work

The primary objective of the current work is to provide a workaround for the previous issue by deriving a robust version of the SURE-criterion that allows for effective dimension estimation under the elliptical model (1) in the presence of heavy-tailed distributions. As described earlier, such procedures are highly called for in applications such as finance, where assumptions on finite moments are usually deemed unreasonable. Our robust extension of the SURE-criterion is carried out via a plug-in strategy where the covariance matrix in the Gaussian SURE-criterion is replaced with a suitable *scatter matrix*. Especially popular in the community of robust statistics, scatter matrices are a class of statistical functionals that measure the dispersion/scatter/variation in multivariate data while (usually) being far more robust to the impact of heavy tails and outliers than the covariance matrix, see Sect. 3 for their definition and several examples. We consider three different plug-in estimators, depending in which form of the Gaussian SURE-criterion the scatter matrix is plugged in. The first two options lead to analytically simple estimators that depend on the data only through the eigenvalues of the used scatter matrix, much like the classical estimators of dimension (see below). Whereas, the third strategy is more elaborate and involves computing particular derivatives of the scatter functional (and the companion location functional).

As our secondary objective, we conduct an extensive simulation study where the proposed methods are compared to each other and to several (families of) competing estimators from the literature. These competitors include: (i) The classical estimator based on successive asymptotic hypothesis testing for the equality of the final eigenvalues of a chosen scatter matrix (testing for *subsphericity*) (Nordhausen et al. 2021) and its high-dimensional version (Schott 2006). (ii) Variation of the previous estimator where the null distributions are bootstrapped instead of relying on asymptotic approximations (Nordhausen et al. 2021). (iii) The general-purpose procedure for inferring the rank of a matrix from its sample estimate known as the *ladle*, which we apply to select scatter matrices, see Luo and Li (2016). (iv) The SURE-estimator of Ulfarsson and Solo (2015) which can be seen as the non-robust version of our proposed estimator. (v) The information criteria-type estimators proposed in Wax and Kailath (1985). (vi) The Bayesian estimator in Minka (2000) that is based on choosing the dimension which yields maximal approximate probability of observing the current data set. We are not aware of comparisons of similar magnitude being conducted earlier in the literature.

Note that further examples of dimension estimators for PCA naturally exist in the literature. For example, Gai et al. (2008), Gai and Stevenson (2010) use a Bayesian approach under a heavy-tailed model, Deng and Craiu (2023) rely on penalized likelihood with aggregation over the tuning parameter values and Zhao et al. (1986) use

information theoretic criteria in a complex-valued version of the Gaussian factor model (2). However, in Sect. 5, we have limited our comparisons to the list of methods given in the preceding paragraph, in particular due to their flexibility in choosing the used scatter matrix, see Sect. 5 for further details.

To summarize the results of our simulation study (given in Sect. 5), they reveal that the SURE-based robust methodology for the determination of the latent dimension is: (i) *Accurate*, achieving good estimation results in various data scenarios. (ii) *Flexible*, that is, it allows the free selection of the used robust scatter matrix. This is in strict contrast to its closest competitor, the asymptotic hypothesis test mentioned above, which is (for theoretical reasons) “locked” to operate with a specific, slow-to-compute scatter matrix. (3.) *Fast*, requiring no bootstrap replicates or any kind of resampling.

1.5 Organization of the manuscript

The manuscript is organized as follows. In Sect. 2 we recall the Gaussian SURE-criterion of Ulfarsson and Solo (2015) for estimating the latent dimension. In Sects. 3 and 4 we propose three different robust extensions of the criterion through the use of different pairs of location and scatter functionals. Sections 5 and 6 contain the simulation study and an empirical (financial) example on asset returns, respectively. In Sect. 7 we conclude with some future research ideas. The proofs of all technical results are collected in Appendix A.

2 SURE criterion for Gaussian PCA

In this section, we recall how the SURE-criterion can be used to estimate the latent dimension d under the Gaussian model (2). Our derivation of the criterion differs from the original version (Ulfarsson and Solo 2015) in that we employ empirical centering of the data, whereas Ulfarsson and Solo (2015) did not. We made this change to the method as it is unreasonable to assume that the true location of the data is known in practice.

Due to the empirical centering, we assume, without loss of generality, that $\mu = 0$ throughout the following. As in Ulfarsson and Solo (2015), we use Stein’s Lemma to construct an unbiased estimator of the risk associated with estimating the signals $V_0 y_i$ by their reconstructions \hat{x}_i based on the first k principal components. Assuming that $k = 1, \dots, p$ is fixed from now on and letting $U_k \in \mathbb{R}^{p \times k}$ denote a matrix of (any) first k orthogonal eigenvectors of the covariance matrix $S_0 := (1/n) \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})'$, the reconstructions can be written as $\hat{x}_i \equiv \hat{x}_i(k) = t_0 + P_k(x_i - t_0)$ where $P_k := U_k U_k'$ is the orthogonal projection onto the space spanned by the first k eigenvectors of S_0 and $t_0 := (1/n) \sum_{i=1}^n x_i$ is the mean vector. For convenience, we replicate an intermediate result towards the final Gaussian SURE-criterion below as Lemma 1. In the lemma the reconstructions \hat{x}_i are treated as functions of the original data x_1, \dots, x_n and the result implicitly assumes the former to be differentiable in the latter, sufficient conditions for which will be discussed later in Sect. 3. In Lemma 1, and throughout the paper, $\|\cdot\|$ denotes the Euclidean norm.

Lemma 1 Under model (2), the quantity

$$R_{1k} := \text{tr}\{(I_p - P_k)S_0\} + \frac{2\sigma^2}{n} \sum_{i=1}^n \sum_{j=1}^p \frac{\partial}{\partial x_{ij}} \hat{x}_{ij} - p\sigma^2$$

is an unbiased estimator of the risk $(1/n) \sum_{i=1}^n \mathbb{E}\|\hat{x}_i - V_0 y_i\|^2$.

The two k -dependent terms of R_{1k} in Lemma 1 have natural interpretations: The term $\text{tr}\{(I_p - P_k)S_0\}$ measures the total variation of the data in directions orthogonal to the first k eigenvectors and takes large values when the used number of eigenvectors is insufficient to capture the full d -variate latent signal. The quantity $(1/n) \sum_{i=1}^n \sum_{j=1}^p (\partial/\partial x_{ij}) \hat{x}_{ij}$ measures the average influence an observation has on their own reconstruction and is often interpreted as the generalized degrees of freedom of the model, where large values indicate overfitting to the data set, see Tibshirani and Taylor (2012) (in the extreme case with $d = p$ we actually have $\hat{x}_{ij} = x_{ij}$). Thus, R_{1k} can be seen to be similar in form to Akaike’s information criterion (AIC) (and other related information criteria), whose two terms also measure model fit and model complexity, respectively.

To apply the criterion R_{1k} in practice, we require an expression for the partial derivatives in Lemma 1. As is shown later in the context of Lemma 3 in Sect. 4, the partial derivatives exist under the assumption that the eigenvalues of S_0 are simple (which holds almost surely under the model (2)), and have the forms shown in Lemma 2 below. Its proof is omitted as the result is a direct consequence of Lemma 1 and Lemmas 3, 4 in Sect. 4. See Ulfarsson and Solo (2008, 2015) for similar results.

Lemma 2 Under model (2), the quantity

$$R_{2k} := \text{tr}\{(I_p - P_k)S_0\} + \frac{2\sigma^2}{n} \sum_{j=1}^k \sum_{\ell=k+1}^p \frac{s_j + s_\ell}{s_j - s_\ell} + \frac{\sigma^2}{n} \{2p + 2(n - 1)k - np\},$$

where $s_1 > \dots > s_p$ are the eigenvalues of S_0 , is an unbiased estimator of the risk $(1/n) \sum_{i=1}^n \mathbb{E}\|\hat{x}_i - V_0 y_i\|^2$.

To apply the criterion R_{2k} in practice, an estimator for the unknown error variance σ^2 is needed and several feasible alternatives exist. For example, Luo and Li (2021) used, in a similar context, the median of the smallest $\lfloor p/2 \rfloor$ eigenvalues of S_0 . The resulting estimator is accurate but makes the implicit assumption that $d \leq \lfloor p/2 \rfloor$. To avoid such difficult-to-verify conditions, we prefer to instead use the final eigenvalue s_p of S_0 as the estimator of the noise variance, imposing minimal assumptions on the latent dimensionality (i.e., that $d < p$). Naturally, the price to pay is that s_p suffers from underestimation in finite samples. Note that, to combat the underestimation, Ulfarsson and Solo (2008) proposed an alternative estimator of σ^2 based on the limiting spectral distribution of the covariance matrix under high-dimensional Gaussian data. Mimicking this strategy is not viable in our scenario as any results on the limiting spectral distributions of the scatter matrices used in Sect. 3 are still scarce in the literature. See also Sect. 7 for further discussion of this choice.

Plugging in the estimator s_p and observing that $\text{tr}\{(I_p - P_k)S_0\} = \sum_{\ell=k+1}^p s_\ell$ now leads to two different sample forms for the SURE criterion for Gaussian PCA:

$$\begin{aligned} \hat{R}_{1k} &:= \sum_{\ell=k+1}^p s_\ell + \frac{2s_p}{n} \sum_{i=1}^n \sum_{j=1}^p \frac{\partial}{\partial x_{ij}} \hat{x}_{ij} - ps_p, \\ \hat{R}_{2k} &:= \sum_{\ell=k+1}^p s_\ell + \frac{2s_p}{n} \sum_{j=1}^k \sum_{\ell=k+1}^p \frac{s_j + s_\ell}{s_j - s_\ell} + \frac{s_p}{n} \{2p + 2(n - 1)k - np\}. \end{aligned} \tag{4}$$

The “hat” notation for \hat{R}_{1k} , \hat{R}_{2k} signifies the fact that they have been obtained from R_{1k} , R_{2k} by replacing the unknown σ^2 with its estimator s_p . In the following sections we obtain outlier-resistant alternatives to both \hat{R}_{1k} and \hat{R}_{2k} via plugging in robust measures of location and scatter in place of the mean vector and covariance matrix in (4). In addition, we will also consider an “asymptotic” version of the criterion \hat{R}_{2k} ,

$$\hat{R}_{3k} := \sum_{\ell=k+1}^p s_\ell + s_p(2k - p), \tag{5}$$

where the terms of the order $o_p(1)$ (in the asymptotic regime where $n \rightarrow \infty$) have been removed. Note that even though we might have $s_j - s_\ell \rightarrow_p 0$ for some indices j, ℓ , the limiting distribution of $\sqrt{n}(s_j - s_\ell)$ for such indices is absolutely continuous (with respect to the Lebesgue measure) (Anderson 1963), meaning that the impact of the double sum in \hat{R}_{2k} can be expected to be negligible for large n .

3 Robust plug-in SURE criteria

Plug-in-techniques are a typical way to create outlier-resistant versions of standard multivariate methods in the community of robust statistics, see, for example, Croux and Haesbroeck (2000), Nordhausen and Tyler (2015), Fan et al. (2021). In this spirit, we replace the mean vector t_0 and the covariance matrix S_0 in the SURE criteria (4), (5) with a pair (t, S) of *location and scatter functionals* (Oja 2010), the definitions of which we recall next. Letting F be an arbitrary p -variate distribution, a location functional (location vector) t is a map $F \mapsto t(F) \in \mathbb{R}^p$ such that, for any invertible $A \in \mathbb{R}^{p \times p}$ and $b \in \mathbb{R}^p$, we have $t(F_{A,b}) = At(F) + b$ where $F_{A,b}$ is the distribution of the random vector $Ax + b$ and $x \sim F$. Similarly, a scatter functional (scatter matrix) S is a map taking values in the space of positive semi-definite matrices and obeying, for any invertible $A \in \mathbb{R}^{p \times p}$ and $b \in \mathbb{R}^p$, the transformation rule $S(F_{A,b}) = AS(F)A'$. These transformation properties are typically referred to as *affine equivariance*.

Location and scatter functionals mimic the properties of the mean vector and the covariance matrix and typically measure some aspects of the center and spread of a distribution, respectively. In particular, if F is the elliptical model (1), then $t(F) = \mu$ and $S(F) = \tau_{S,F}VD^2V'$ for all location and scatter functionals (t, S) for which $t(F)$ and $S(F)$ exist, where the scalar $\tau_{S,F} > 0$ depends on both the exact distribution

of the spherical z_i and on the used scatter functional, see [Oja (2010),Theorem 3.1]. Hence, all choices of (t, S) estimate, up to scale, the same quantities under the elliptical model, implying that the replacing of the mean vector and the covariance matrix in SURE with the pair (t, S) is warranted (at least in the Gaussian special case (2) of the elliptical model, under which the SURE criteria in Sect. 2 were derived). Note that this equivalence of different (t, S) -pairs under elliptical data does not necessarily mean that the sample dimension estimates given by different choices of (t, S) should always be equal. Namely, the equivalence indeed holds under the population level model (1), but in practical situations the accuracy of the estimates is greatly influenced by the finite-sample properties (in particular, robustness properties) of the used location and scatter functionals. This is clearly evident in the simulation results of Sect. 5.

Examples of popular location and scatter functionals are given later in this section and we assume, for now, that we have selected some robust location-scatter pair (t, S) . Outlier-resistant versions of the forms \hat{R}_{2k} and \hat{R}_{3k} of the Gaussian SURE criterion in (4) and (5) are then straightforwardly obtained. Namely, we simply replace the eigenvalues s_j of the covariance matrix S_0 with the eigenvalues of the scatter functional S in the definitions. Note that while the location functional t does not play an explicit role in this construction, it is usually a part of the definition of S , see for example the spatial median and the spatial sign covariance matrix later in this section.

As an alternative to the above, rather simplistic plug-in estimators, we consider also a more elaborate extension based on the form \hat{R}_{1k} of SURE in (4) where, in addition to S_0 , we also replace the partial derivatives $(\partial/\partial x_{ij})\hat{x}_{ij}$ with their counterparts based on the robust pair (t, S) . That is, the robust version of \hat{R}_{1k} uses the reconstruction estimates $\hat{x}_i = t + P_k(x_i - t)$ where the centering is done with the robust location functional t (instead of t_0) and the projection matrix P_k is now taken to be onto the space spanned by the first k eigenvectors of the robust scatter functional S (instead of S_0). Due to its more technical nature in comparison to the other two criteria, we have postponed the discussion of the extension of \hat{R}_{1k} to Sect. 4.

Finally, regardless of which of the three criteria \hat{R}_{1k} , \hat{R}_{2k} and \hat{R}_{3k} one uses, the corresponding estimate \hat{d} of the latent dimension d is obtained as the minimizing index,

$$\hat{d} = \operatorname{argmin}_{k=0,\dots,p-1} \hat{R}_{jk}.$$

We next recall several popular options for the location-scatter pair (t, S) .

3.1 Mean vector and covariance matrix

The most typical choice for the pair (t, S) is the mean vector and the covariance matrix, i.e.,

$$t(F) = E_F(x), \quad S(F) = E_F[\{x - t(F)\}\{x - t(F)\}'], \tag{6}$$

where E_F means that the expectation is taken under the assumption that $x \sim F$. This choice simply leads to the Gaussian SURE-criterion discussed in Sect. 2. As discussed

before, this option is, despite often being the optimal choice under the assumption of normality, also highly non-tolerant against outliers and heavy tails.

3.2 Spatial median and spatial sign covariance matrix

The spatial median $t(F)$ of a distribution F is defined as any minimizer of the convex function

$$t \mapsto E_F\{\|x - t\| - \|x\|\},$$

over $t \in \mathbb{R}^p$. The spatial median is one of the oldest and most studied robust measures of multivariate location, see, Haldane (1948), Brown (1983), and reverts to the univariate concept of median when $p = 1$. It can be shown to exist for any F (in particular, no moment conditions are required) and it is unique as soon as F is not concentrated on a line in \mathbb{R}^p (Milasevic and Ducharme 1987) which is guaranteed, in particular, almost surely when F is absolutely continuous.

The standard scatter functional counterpart for the spatial median is the spatial sign covariance matrix (SSCM), defined as,

$$S(F) = E_F \left\{ u(x - t(F))u(x - t(F))' \right\},$$

where $t(F)$ is the spatial median of F , which is assumed to be unique, and the sign function $u : \mathbb{R}^p \rightarrow \mathbb{R}^p$ is defined as $u(x) = x/\|x\|$ for $x \neq 0$ and $u(0) = 0$. Like its location counterpart, also the SSCM has been extensively studied in the literature, see, for instance, Marden (1999), Visuri et al. (2000), Dürre et al. (2016), Bernard and Verdebout (2021).

The defining feature of SSCM is that it depends on the data only through the “signs” $u(x - t(F))$, giving equal weight to points in a given direction regardless of their norm (which, in turn, is what makes the SSCM robust to outliers). Especially for high-dimensional data, this loss of information introduced by the discarding of the observation magnitudes is relatively small as it represents losing only a single degree of freedom in the p -dimensional space (whereas the sign contains the remaining $p - 1$ degrees of freedom).

We also note that the spatial median and the spatial sign covariance matrix are, strictly speaking, not a pair of location and scatter functionals in the usual sense as they satisfy the equivariance properties listed in Sect. 2 only when the matrix A is orthogonal. However, this is not an issue in our scenario for the following reasons: (i) The spatial median is a consistent estimator of the location parameter in the elliptical model (1) under minor regularity conditions, see, e.g., Magyar and Tyler (2011). (ii) Under the elliptical model (1), two eigenvalues s_j, s_ℓ of the SSCM are equal if and only if the corresponding elements σ_j, σ_ℓ of the matrix D are equal, see Dürre et al. (2016). Hence, the eigenvalues of the spatial sign covariance matrix contain the same (qualitative) information about the latent signal dimension as those of any “proper”, affine equivariant scatter functional. However, the spatial sign covariance matrix is also known to non-linearly compress the range of the eigenvalues (Vogel and Fried

2015), making it more difficult to distinguish between the signal and the noise and, thus, we next consider an alternative to it. This alternative is known as Tyler’s shape matrix and is often seen as the affine equivariant version of the SSCM.

3.3 Tyler’s shape matrix

Tyler’s shape matrix (Tyler 1987) is one of the earliest proposed and most studied scatter functionals, see, e.g., Dümbgen and Tyler (2005), Wiesel (2012). Using it requires a location functional t and, in the following, we take this to be the spatial median, as is common in the literature. Tyler’s shape matrix $S(F)$ is defined as any S with $\det(S) = 1$ and satisfying the following fixed-point equation,

$$E_F \left[u \left(S^{-1/2} \{x - t(F)\} \right) u \left(S^{-1/2} \{x - t(F)\} \right)' \right] = \frac{1}{p} I_p. \tag{7}$$

A unique solution $S(F)$ is obtained as soon as F does not concentrate too heavily on a subspace in \mathbb{R}^p , see Dümbgen and Tyler (2005) for the exact conditions. Inspection of the Eq. (7) also reveals that any solution S to it is defined only up to its scale and, to obtain a unique representative, a popular choice is indeed to use the determinant condition $\det(S) = 1$ to fix the scale of solution, see Paindaveine (2008). Consequently, $S(F)$ does not describe the full scatter of F but only its *shape* (scale-standardized scatter). However, this is sufficient for our purposes as scaling preserves the ordering of the eigenvalues of $S(F)$ and, hence, their division into signal and noise. Note that this also means that Tyler’s shape matrix satisfies the affine equivariance property discussed in the beginning of Sect. 3 only up to scale, $S(F_{A,b}) = \{\det(A)\}^{-2/p} AS(F)A'$.

The computation of $S(F)$ can be shown to correspond to a geodesically convex minimization problem (Wiesel 2012), meaning that an efficient algorithm for its estimation in practice is straightforwardly constructed. In our simulations we have used the R-package ICSNP (Nordhausen et al. 2018) for this purpose.

3.4 Hettmansperger-Randles estimator

As our final choice for the location scatter pair (t, S) , we consider the so-called Hettmansperger–Randles (H–R) estimator, which was originally introduced in the context of robust location estimation (and the associated shape functional was obtained as a “by-product” of the location estimation) (Hettmansperger and Randles 2002). The H–R pair $(t(F), S(F))$ is defined as any (t, S) , with $\det(S) = 1$ and satisfying the following pair of fixed-point equations,

$$\begin{aligned} E_F \left\{ u(S^{-1/2}(x - t)) \right\} &= 0 \\ E_F \left\{ u(S^{-1/2}(x - t))u(S^{-1/2}(x - t))' \right\} &= \frac{1}{p} I_p. \end{aligned} \tag{8}$$

Observing that the LHS of the first Eq. in (8) is, disregarding the matrix $S^{-1/2}$, the gradient of the objective function (6) of the spatial median, we see that the H–R pair $(t(F), S(F))$ can be interpreted as simultaneously determined spatial median and Tyler’s shape matrix. This concurrent estimation of location and scatter (or, rather, shape as again any solution S to the fixed-point equations is unique at most up to scale) then makes the resulting estimator affine equivariant (up to scale in case of S).

Despite its attractiveness, the theoretical properties of the H–R estimator have garnered less attention in the literature when compared to its previously introduced alternatives. In particular, we are not aware of any studies investigating conditions that would guarantee the uniqueness of the solution (t, S) .

To summarize, the three robust alternatives to the mean-covariance pair introduced in this section can all be seen to estimate analogous quantities, while at the same time forming a sort of “hierarchy” with respect to their equivariance properties: (i) the spatial median and SSCM satisfy affine equivariance only for orthogonal A , (ii) replacing SSCM with Tyler’s shape matrix yields the full affine equivariance property for the scatter (shape) functional and, (iii) both the location and scatter (shape) components of the H–R estimate are affine equivariant. As affine equivariance is the natural transformation property for a scatter functional to have in the presence of the elliptical model (1), we thus expect that the previous ordering applies also to the corresponding SURE-procedures’ comparative performances in practice. This claim will be investigated through simulations in Sect. 5.

4 Robust extension of \hat{R}_{1k}

In this section, we explore extending the SURE-criterion \hat{R}_{1k} in (4) to accommodate an arbitrary location-scatter pair. The theoretical cost of such an extension is considerably larger than for \hat{R}_{2k} and \hat{R}_{3k} as, instead of simply plugging in eigenvalues, it involves computing the partial derivatives $(\partial/\partial x_{ij})\hat{x}_{ij}$.

In the sequel, let the observed sample x_1, \dots, x_n of points in \mathbb{R}^p be fixed and denote its empirical distribution by F_n . Moreover, for $\varepsilon > 0$ we let $F_{n,i,j,\varepsilon}$ denote the empirical distribution of the perturbed sample $x_1, \dots, x_i + \varepsilon e_j, \dots, x_n$ where e_j is the j th vector in the canonical basis of \mathbb{R}^p . For the extension \hat{R}_{1k} to be well-defined in the first place, t and S are naturally required to be differentiable in a suitable sense, and the next assumption formalizes this requirement.

Assumption 1 For any $i = 1, \dots, n$ and $j = 1, \dots, p$, there exists $h_{ij} \in \mathbb{R}^p$ and a symmetric $H_{ij} \in \mathbb{R}^{p \times p}$ satisfying

$$\frac{1}{\varepsilon}\{t(F_{n,i,j,\varepsilon}) - t(F_n)\} \rightarrow h_{ij} \quad \text{and} \quad \frac{1}{\varepsilon}\{S(F_{n,i,j,\varepsilon}) - S(F_n)\} \rightarrow H_{ij},$$

as $\varepsilon \rightarrow 0$.

In order for also the projection matrix P_k (onto the span of the first k eigenvectors of $S(F_n)$) to be differentiable in the previous sense for all $k = 1, \dots, p$, all eigenvalues of the matrix $S(F_n)$ must be simple. This condition, formalized in Assumption 2

below, is rather mild and, in particular, holds almost surely for both the covariance matrix and the SSCM if the points x_1, \dots, x_n are drawn from an absolutely continuous distribution.

Assumption 2 The eigenvalues of $S(F_n)$ are distinct.

Under the previous two assumptions, the partial derivatives $(\partial/\partial x_{ij})\hat{x}_{ij}$ exist and their sum over j has the analytical form given in the next lemma.

Lemma 3 Under Assumptions 1 and 2, we have

$$\sum_{j=1}^p \frac{\partial}{\partial x_{ij}} \hat{x}_{ij} = k + \sum_{j=1}^p \text{tr}\{(I_p - P_k)h_{ij}e'_j\} + \sum_{j=1}^p e'_j A_{ij}\{x_i - t(F_n)\},$$

where

$$A_{ij} := \sum_{\ell=1}^k \sum_{m=k+1}^p \frac{1}{s_\ell - s_m} (T_\ell H_{ij} T_m + T_m H_{ij} T_\ell),$$

T_ℓ is the orthogonal projection onto the space spanned by the ℓ th eigenvector of $S(F_n)$, and s_ℓ is the corresponding eigenvalue.

Lemma 3 essentially says that, as soon as one obtains expressions for the quantities h_{ij} and H_{ij} in Assumption 1 (and Assumption 2 holds) for some particular location scatter pair (t, S) , these can be plugged in to Lemma 3 to construct a version of the SURE-criterion \hat{R}_{1k} that is based on (t, S) . In Lemma 4 below we have provided, for completeness, these expressions for the standard mean-covariance pair. The resulting SURE-criterion \hat{R}_{1k} is, naturally, the Gaussian SURE as described in Sect. 2.

Lemma 4 The mean vector and the covariance matrix satisfy Assumption 1 with

$$h_{ij} = \frac{1}{n}e_j, \quad H_{ij} = \frac{1}{n}e_j\{x_i - t(F_n)\}' + \frac{1}{n}\{x_i - t(F_n)\}e'_j.$$

Despite not offering us anything new, Lemma 4 also serves in its simplicity as a contrast to our next result, detailing the forms of h_{ij} and H_{ij} for the spatial median/SSCM-pair. What makes deriving these quantities more complicated, compared to the mean-covariance pair, is the fact that no analytical expression is available for the spatial median (instead, it is obtained as a minimizer of the objective function described in Sect. 3).

Lemma 5 Assume (i) that the points x_1, \dots, x_n are not concentrated on a line in \mathbb{R}^p and, (ii) that $t(F_n) \neq x_i$, for all $i = 1, \dots, n$. Then the spatial median and the spatial sign covariance matrix satisfy Assumption 1 with

$$h_{ij} = G^{-1} A_i e_j,$$

$$H_{ij} = \frac{1}{n} \left\{ A_i e_j \frac{y'_i}{\|y_i\|} + \frac{y_i}{\|y_i\|} e'_j A_i - \sum_{\ell=1}^n A_\ell G^{-1} A_i e_j \frac{y'_\ell}{\|y_\ell\|} - \sum_{\ell=1}^n \frac{y_\ell}{\|y_\ell\|} e'_j A_i G^{-1} A_\ell \right\},$$

where $A_i := w_i(I_p - y_i y'_i / \|y_i\|^2)$, $w_i := \|y_i\|^{-1}$, $y_i := x_i - t(F_n)$ and $G := \sum_{i=1}^n A_i$.

Plugging h_{ij} and H_{ij} from Lemma 5 to the derivatives in Lemma 3 and consequently to \hat{R}_{1k} in (4) now gives us yet another robust criterion for determining the signal dimension. We note that while the additional assumption (ii) imposed in Lemma 5 seems to be difficult to analyze theoretically, its validity is nevertheless simply checked in practice (and the assumption (i) is satisfied, in particular, almost surely when F is absolutely continuous).

Mimicking the proof of Lemma 5 it would next be possible to derive equivalent results also for our two remaining location-scatter pairs. However, we have decided not to do so and the reasons for this are two-fold: (i) Some preliminary computations (not shown here) show that these computations lead, as with the spatial median/SSCM-pair in Lemma 5, to analytically cumbersome expressions for h_{ij} and H_{ij} , from which no real insight can be gained. (ii) Due to the complexity of the resulting expression (and the large number of nested summations involved), the practical usefulness of the extensions is questionable. Indeed, as our timing comparisons in Sect. 5.4 demonstrate, the version of \hat{R}_{1k} obtained based on Lemma 5 is several orders of magnitude inferior to \hat{R}_{2k} and \hat{R}_{3k} in computational speed, while at the same time offering no or only minuscule gains in accuracy. Some preliminary exploration reveals that this issue is still further magnified for Tyler’s shape matrix. Hence, while these extensions would be technically possible to derive, we did not see any real practical value in doing so.

5 Simulations

In order to study the finite-sample properties of our proposed robust extensions of SURE, we conduct an array of simulation studies. As competing methods we have used the following set of well-established estimators from the literature, see Sect. 1.4 for further details. Naturally, also other choices would be available, but this particular set was chosen as (i) it contains representatives of estimators based both on asymptotic results and on computationally intensive ideas, and (ii) all methods in the set allow choosing the used scatter matrix (at least to some extent), letting us separately compare the different methodologies and the different levels of robustness.

- (i) The classical estimator based on an asymptotic test of subsphericity (Schott 2006; Nordhausen et al. 2021). The R-package `ICtest` (Nordhausen et al. 2021) includes two implementations of it, one based on the covariance matrix and one based on the H–R estimator, and we include both of them in the comparison. Additionally, we include the high-dimensional variant of the test (Schott 2006) which is based on the covariance matrix.
- (ii) The same estimator as (i) but with the null distribution of the test estimated through bootstrap. This estimator can be based on any of the four scatter matrices described

- in Sect. 3 and we thus include all of them in the comparison. We used 200 bootstrap samples throughout the study, the default value in the implementation in `ICtest` (Nordhausen et al. 2021).
- (iii) The ladle estimator of Luo and Li (2016) which, too, can be based on any of the four scatter matrices. The estimator is based on resampling, for which we used the default value 200 in the implementation in `ICtest` (Nordhausen et al. 2021).
 - (iv) Our centered version of the SURE-estimator of Ulfarsson and Solo (2015).
 - (v) The AIC and MDL criteria from Wax and Kailath (1985), see their Eqs. (16) and (17). We use both criteria separately with the covariance matrix and the H–R estimator. The derivation in Wax and Kailath (1985) does not cover the latter case (H–R), meaning that it can be considered as an “experimental” robust plug-in version of their method.
 - (vi) The Bayesian approach in Minka (2000), see their Eq. (30). As in item (v) above, we apply the method both with the covariance matrix and the H–R estimator, making also this estimator experimental in nature.

The above six categories of estimators are denoted in the following as Asymp, Boot, Ladle, SURE, Wax and Minka, respectively, with the used scatter matrix given in parenthesis. E.g., Asymp(HR) denotes the asymptotic test based method using the H–R estimator. Additionally, we denote the high-dimensional variant of Asymp by Schott. In addition, we distinguish three different versions of the SURE-estimator, SURE1, SURE2 and SURE3, referring to using the objective functions \hat{R}_{1k} , \hat{R}_{2k} and \hat{R}_{3k} , respectively. We thus have a total of 26 methods to compare, and these have been summarized in Table 1. The final four columns of the table are related to the timing study in Sect. 5.4. The R-implementations of the methods are available at <https://users.utu.fi/jomivi/software/>.

5.1 Tail thickness

In the first simulation study, we explore how the methods perform under varying levels of heavy-tailedness. As a setting for this, we consider multivariate t -distributions with the degrees of freedom equal to $\nu = 1, 3, 5, \dots, 25$. Thus, the heaviest tails are obtained in the case $\nu = 1$, corresponding to the multivariate Cauchy distribution. The simulation is repeated 100 times for every degree of freedom, and for every repetition a random sample consisting of $n = 100$ observations is generated. In each case, we take the latent dimension to be $d = 6$, whereas as the total dimensionality we use $p = 10$. The error “variance” (i.e., the square of the final diagonal elements of D in (1)) is always $\sigma^2 = 0.5$ and the signal “variances” (i.e., the squares of the first d diagonal elements of D in (1)) are randomly generated from the uniform distribution $\text{Unif}(1, 3)$, independently for each of the 100 repetitions. The proportions of correctly estimated dimensions d for each of the 26 methods are presented in Fig. 1, divided into two panels for visual convenience.

Unsurprisingly, the classical covariance matrix based methods fail to consistently find the correct latent dimension d in the presence of too heavy tails (left sides of the plots). This effect is most pronounced in the case $\nu = 1$ where the corresponding t -distribution does not possess the finite second-order moments required by the

Table 1 The estimators included in the simulation study, along with binary indicator for their robustness

Method (scatter matrix)	Is robust?	$n = 200$		$n = 400$	
		$p = 10$	$p = 20$	$p = 10$	$p = 20$
SURE1(Cov)	No	0.01	0.00	0.00	0.00
SURE1(SSCM)	Yes	0.13	2.04	0.19	2.74
SURE2(SSCM)	Yes	0.01	0.02	0.01	0.02
SURE2(Tyler)	Yes	0.02	0.02	0.02	0.03
SURE2(HR)	Yes	0.06	0.08	0.10	0.15
SURE3(Cov)	No	0.00	0.00	0.00	0.00
SURE3(SSCM)	Yes	0.01	0.01	0.01	0.02
SURE3(Tyler)	Yes	0.01	0.02	0.02	0.03
SURE3(HR)	Yes	0.07	0.09	0.09	0.15
Asymp(Cov)	No	0.00	0.00	0.00	0.00
Asymp(HR)	Yes	0.06	0.08	0.09	0.12
Schott(Cov)	No	0.01	0.00	0.00	0.01
Boot(Cov)	No	0.23	0.32	0.46	1.50
Boot(SSCM)	Yes	2.32	5.50	4.17	10.77
Boot(Tyler)	Yes	3.04	7.67	5.01	13.16
Boot(HR)	Yes	9.63	24.97	13.79	38.65
Ladle(Cov)	No	0.09	0.15	0.09	0.18
Ladle(SSCM)	Yes	2.01	2.21	3.26	3.88
Ladle(Tyler)	Yes	3.17	4.39	4.60	6.29
Ladle(HR)	Yes	18.35	27.69	19.19	28.65
Wax AIC(Cov)	No	0.00	0.00	0.00	0.00
Wax AIC(HR)	Yes	0.06	0.08	0.09	0.13
Wax MDL(Cov)	No	0.00	0.00	0.00	0.00
Wax MDL(HR)	Yes	0.06	0.09	0.09	0.12
Minka(Cov)	No	0.00	0.00	0.00	0.00
Minka(HR)	Yes	0.06	0.09	0.09	0.13

The final four columns give the average running times (in seconds) of the methods over 10 replicates in various settings

covariance estimation. The robust methods, on the other hand, do not suffer from this issue as they make no moment assumptions on the data generating distribution. As t -distributions with low degrees of freedom regularly produce observations that would be classified as outliers in the standard (Gaussian) statistical practice, the corresponding simulation settings can also be interpreted to measure how well the methods perform in the presence of a large number of outlying observations.

As the degrees of freedom increase, we observe that the covariance based methods start to outperform the robust alternatives. The reason for this is that, when $\nu \rightarrow \infty$, the multivariate t -distribution approaches the normal distribution for which the covariance based methods offer optimal inference.

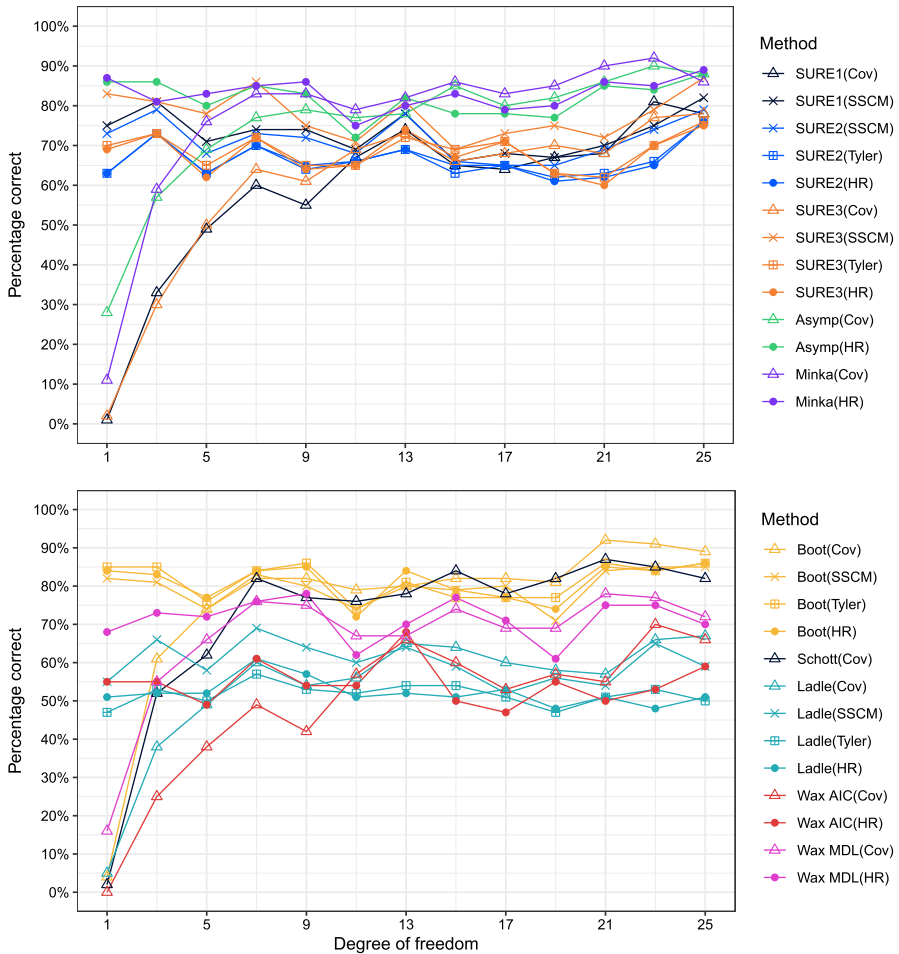


Fig. 1 Percentage of correctly estimated dimensions d as a function of the degree of freedom ν of the multivariate t -distribution in the tail thickness simulation. The sample size is $n = 100$ throughout

As the performance criterion in Fig. 1, proportion of correct estimates, gives no indication of any possible under/overestimation, we show in Fig. 2 the mean estimation errors for each combination of setting and method. Thus, values close to zero in Fig. 2 indicate unbiased estimation. We observe that several of the methods consistently overestimate the true dimension, which is actually desirable over underestimation as it means that no signal information is lost (at the cost of introducing more noise). Boot, Asymp and Minka are the most unbiased estimators, whereas SURE-based methods tend to overestimate by 0.25–0.50 dimensions on average.

Comparing the methods by type (different colours in Figs. 1 and 2), the overall best performances are given by the robust bootstrap (Boot, yellow), asymptotic (Asymp, green) and Bayesian (Minka, purple) methods, with marginal differences between the two. Somewhat surprisingly, the computationally most expensive method, i.e.,

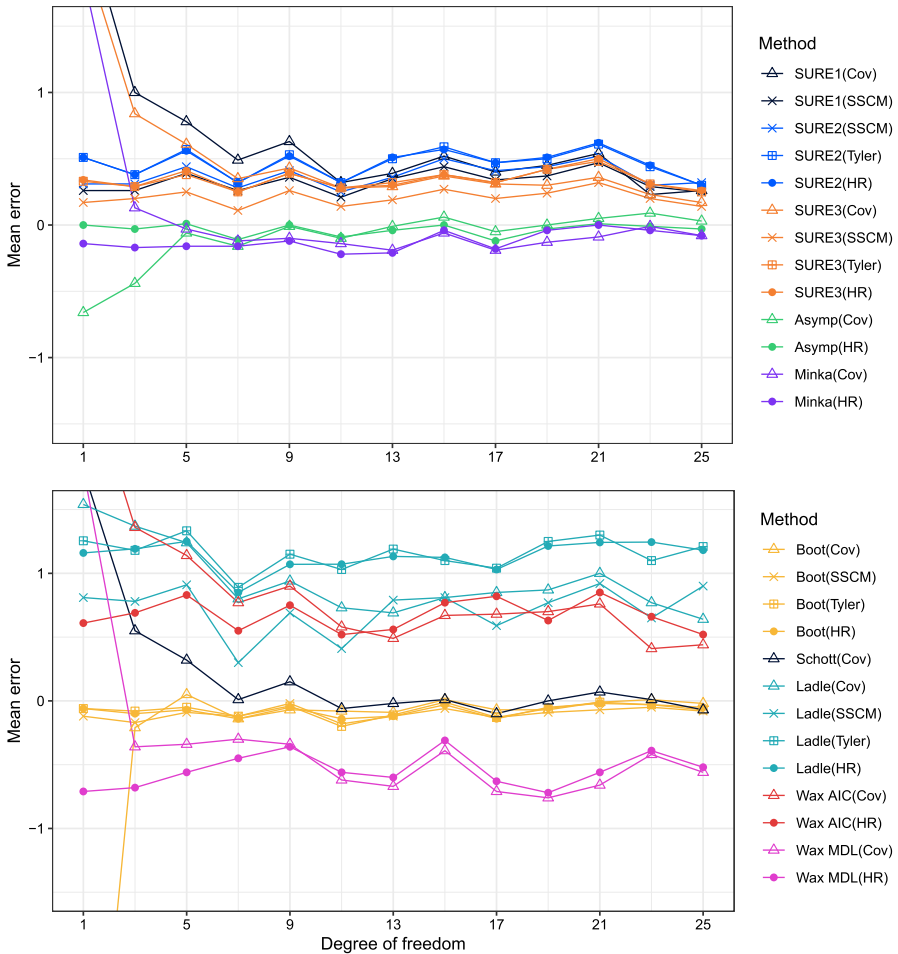


Fig. 2 Mean of the error in dimension estimation as a function of the degree of freedom ν of the multivariate t -distribution in the tail thickness simulation. The sample size is $n = 100$ throughout

the ladle, has a relatively bad performance in this scenario. The information criterion type estimators (Wax, red and magenta) give a solid performance but are not among the overall best choices. Comparing the different types of SURE to each other, it seems that the additional computational and theoretical complexity of SURE1 (black) does not provide additional benefits when compared to SURE2 (blue) and SURE3 (orange), which both have a relatively similar performance, not falling much behind Boot, Asymp and Minka.

5.2 Latent dimension

In our second simulation study, we investigate how the relative size of the underlying latent dimension d affects the methods' performances. As the main selling point

of the SURE-based methods is their light computational load, we have dropped the more computationally intensive methods (Boot, Ladle), as well as the less successful information theoretic methods (Wax), from the comparison, focusing in this (and the following) study on comparing SURE only to its most relevant competitors, Asymp and Minka. We choose SURE2 to be the “representative” of the SURE-family as, based on the first simulation study, both SURE1 and SURE3 had performance similar to it. Thus, the families of methods included in the current simulation study are Asymp, SURE2 and Minka.

Recall that SURE2 estimates the latent dimension as the index minimizing the corresponding objective function $k \mapsto \hat{R}_{2k}$. Based on our experiments, this strategy can sometimes be quite unstable, especially when the latent dimension is comparatively small. Thus, as an experimental alternative we propose estimating d as the change point in the series of differences $\hat{R}_{2(k+1)} - \hat{R}_{2k}$. To understand the motivation for this, consider the following two typical forms for the graph formed by the points (k, \hat{R}_{2k}) : (i) The points (k, \hat{R}_{2k}) form a V-shaped curve around d . In this case, the true dimension is both a minimizer and a location change point of the differences (the differences change sign at the true dimension). (ii) The graph (k, \hat{R}_{2k}) decreases linearly until d and stays roughly constant afterwards. In this case, d is a location change point of the differences but not necessarily a minimizer (it might happen that the minimizer occurs only after d). Thus, in these two (rather idealistic) examples, the experimental change point alternative offers more consistent detection of the dimension than the standard method of seeking the minimizer. We implemented the change point detection as binary segmentation through the function `cpt.meanvar` in the R-package `changepoint` (Killick and Eckley 2014). The resulting method is denoted in the sequel as “SURE2 cp”.

We consider two sample sizes, $n = 100$ and $n = 1000$. For the former, we fix the total dimensionality to $p = 10$ and let the latent dimension vary as $d = 1, 2, \dots, 9$. For the latter, we use $p = 100$ and $d = 5, 10, 15, \dots, 95$. We repeat the simulation 100 times for every combination of parameters, generating in each repetition a random sample from the multivariate t -distribution with 1 degree of freedom. The error variance is fixed to $\sigma^2 = 0.5$ and the signal variances are randomly generated from the uniform distribution $\text{Unif}(1, 3)$, independently in every repetition. To get a finer comparison between the methods, we use the average absolute estimation error as our performance criterion in this (and the following) simulation study. The average absolute errors for the different dimension estimation procedures are presented in the two panels of Fig. 3, separately for $n = 100$ and $n = 1000$.

The top panel of Fig. 3 reveals that SURE2 and SURE2 cp give the best and most consistent performance when $n = 100$. And, even though Asymp and Minka are slightly better for the smallest values of d , their use cannot be recommended in practical scenarios due to the drop in performance for larger d . The difference between the different scatter matrices is quite minor, but the overall best choices are Tyler and HR, which was somewhat expected as SSCM is not a “proper” scatter matrix, see the discussion in Sect. 3.2. When $n = 1000$ (bottom panel of Fig. 3), Asymp and Minka with HR achieve very good performance, as does SURE2 when $d \geq 50$. The subpar performance of SURE2 for small values of d appears to be a sample size issue as

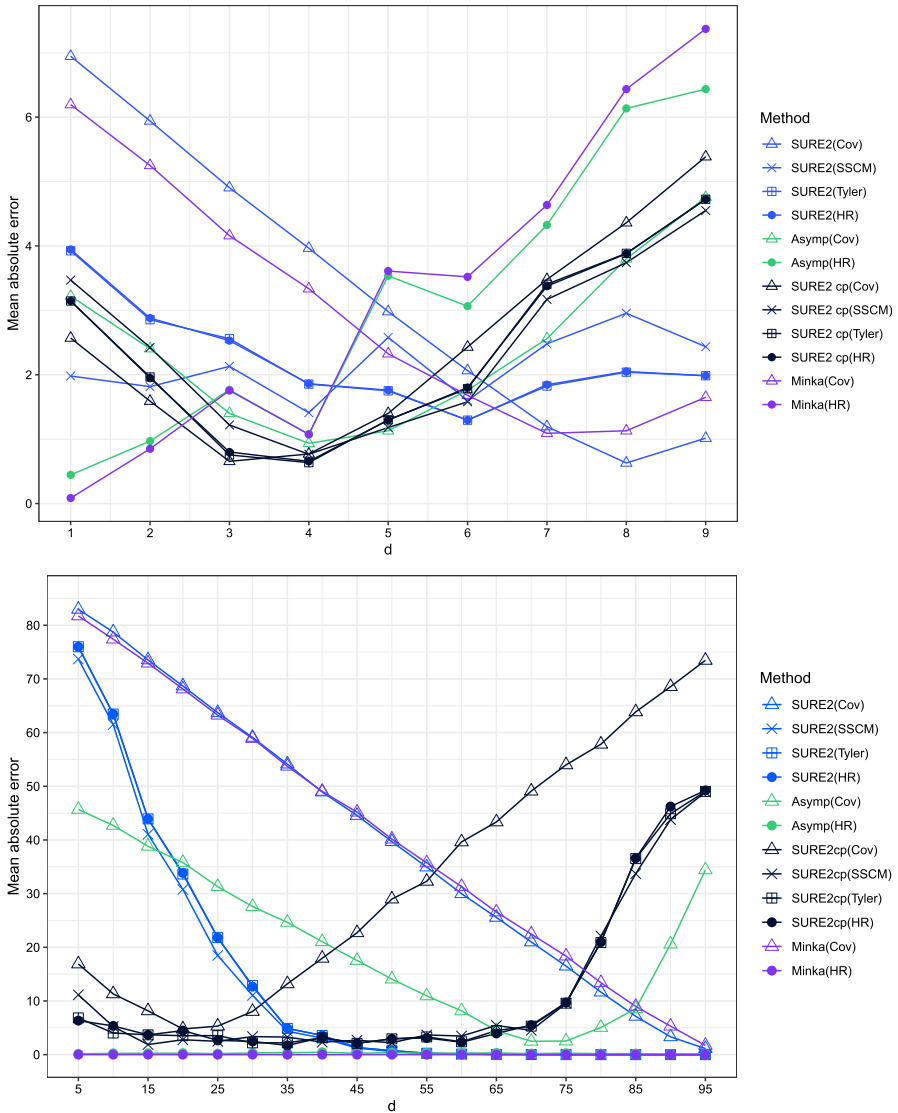


Fig. 3 Average absolute errors of the dimension estimates as a function of the underlying dimension d when sample size is $n = 100$ (top panel) or $n = 1000$ (bottom panel)

when we tried increasing the sample size to $n = 2000$ (not shown here) SURE2 too achieved performance equal to Asymp and Minka. SURE2 cp performs the best when $d \leq 55$, confirming our earlier idea about the usefulness of the changepoint strategy for low latent dimensionalities.

To summarize the results of the study, when n, p are both low, SURE-based methods offer the overall best guarantees for dimension estimation across all values of d ,




Fig. 4 Percentages of correctly estimated dimensions as a function of the sample size n . The scale of the horizontal axis is logarithmic

whereas, when n, p are larger, Minka and Asymp are the most consistent, requiring lower sample sizes than SURE to achieve near-perfect estimation.

5.3 Sample size

In the third simulation, we study the effect of the sample size on the estimation accuracy, including again the same set of methods as in the previous study. The considered sample sizes are $n = 500, 750, 1000, 1500, 2000, 2500, 5000$. The simulation is repeated 100 times for every sample size n , such that for every repetition a random sample of n observations is generated from the multivariate t -distribution with 1 degree of freedom. We take the latent and the total dimensionalities to be $d = 20$ and $p = 100$, respectively, throughout the simulation. As the error variance we use $\sigma^2 = 0.5$ and the signal variances are again randomly generated from the uniform distribution $\text{Unif}(1, 3)$, independently for each of the 100 replicates. The proportions of correctly estimated dimensions d for the different procedures are presented in Fig. 4.

The most striking feature of Fig. 4, which also sheds some light on the behaviour of SURE2 in the previous simulations, is the jump in the mean absolute error of robust SURE2 from 80 to 0 between $n = 750$ and $n = 2000$. Interestingly, the fact that SURE2 cp has more even behaviour across the different sample sizes indicates that this jump is not so much a consequence of the SURE criterion \hat{R}_{2k} itself but of the way in which the dimension estimate is selected based on the criterion values (recall that SURE2 picks the minimizing value of k and SURE2 cp uses a more complicated change point technique). We thus conclude that the standard technique of choosing

simply the minimizing index of the SURE criterion is not optimal unless n is large enough. This matter clearly warrants further investigation and, due to its complexity, we have left it for future research, see Sect. 7. Finally, we also observe that, overall, the used scatter matrix seems to have very little effect on the results, apart from Cov, which again breaks down in the presence of a heavy-tailed distribution.

5.4 Computation time

As our final simulation study, we compare the running times of all 26 methods included in Table 1. The change point variant of SURE2 is not included as its computational difference to the base SURE2-method is marginal and negligible compared to the differences between the methods itself. We distinguish two different sample sizes $n = 200, 400$ and two different dimensionalities $p = 10, 20$, their combinations leading to a total of four different settings. For each setting, we take the data distribution to be the multivariate t -distribution with $\nu = 1$ degrees of freedom, $d = 0.6p$ and the signal and noise variances as in the previous simulation. We run each of the 26 methods 10 times on each setting (using the same set of 10 data for each method) and record their computational times. The experiment was conducted on a desktop computer with AMD Ryzen 5 3600 6-core processor and 16 GB RAM. The average running times in seconds are given in the final four columns of Table 1.

From Table 1 we make the following observations: (i) Computational complexity in the methods stems from two sources, the choice of the scatter matrix and bootstrap replications (as performed by Boot and Ladle), of which the latter has a significantly greater impact on the timing. (ii) The doubling of the sample size n has a quite minor effect on the computational times, whereas the doubling of the dimension p serves to multiply the times by roughly 1.5. (iii) Of the robust methods, the fastest are by far SURE2, SURE3, Asymp, Wax AIC, Wax MDL and Minka which all have computational times roughly of the same order of magnitude. SURE1 falls somewhere in between them and the more intensive Boot and Ladle.

Based on the observations made above and in the previous experiments, we conclude that the SURE-based robust methodology (i) offers a fast and competitive alternative to standard bootstrap-based methods, (ii) partially retains its functionality also at very low sample sizes, unlike Asymp and Minka, and (iii) requires higher sample sizes to reach near perfect estimation results than Asymp and Minka.

6 Application: asset returns

We next illustrate the robust SURE methods in a financial data set that was used to demonstrate principal component analysis in the classical textbook [Tsay (2010), Sect. 9] to search for common latent variables explaining (joint) asset return variability. The data itself is available on the author's (Ruey S. Tsay's) webpage and consists of monthly log stock returns (including dividends) of five stocks (IBM, HPQ, INTC, JPM, BAC)




Fig. 5 The monthly log returns (in percentages) of five stocks (IBM, HPQ, INTC, JPM, BAC) from January 1990 to December 2008. See [Tsay (2010), Section 9]

from January 1990 to December 2008.¹ These $p = 5$ time series of length $T = 228$ months are illustrated in Fig. 5. Tsay (2010) computed the Portmanteau test statistics and found that despite the time series nature of the data there is no substantial serial correlation in returns, and hence we also ignore the serial dependence in our analysis. According to the results of PCA, Tsay (2010) concluded two common latent variables in his interpretation: The “market component” represents the general movement of the stock market and the “industrial component” represents the difference between the two industrial sectors, namely technology (IBM, HPQ and INTC) versus financial services (JPM and BAC). In addition, Tsay (2010) points out that IBM stock “has its own features that are worth further investigation”.

Instead of computing only a single estimate of the latent dimension for the data set, we take a local approach and run a window of length ℓ through the data. For each of the $T - \ell + 1$ windows we then estimate the latent dimension using one of our proposed methods. As the obtained dimensionalities correspond to the windows and not the actual observation months, we “back-transform” them as follows: For each individual month, we take the weighted average of the estimated dimensions of all windows in which that particular month is a member. We assign the weights

¹ See the website <https://faculty.chicagobooth.edu/ruey-s-tsay/research/analysis-of-financial-time-series-3rd-edition>. Specifically, stocks are International Business Machines (IBM), Hewlett-Packard (HPQ), Intel Corporation (INTC), J.P. Morgan Chase (JPM) and Bank of America (BAC).




Fig. 6 The smoothed curves of latent dimensionalities estimated with SURE2 using the window approach as described in the main text. The two panels correspond to the window lengths of $\ell = 48$ and $\ell = 72$ months, respectively

such that the middle two observations in each window get a weight equal to 1 and the weights decrease linearly towards the window endpoints. This procedure thus produces a “smoothed” curve of estimated latent dimensions for the full observation period. To guarantee that the ellipticity of the data is at least partially fulfilled, we resort to rather small window lengths, taking either $\ell = 48$ or $\ell = 72$ in the following. Visual inspection (not shown here) reveals that scatter plots of the obtained windows indeed exhibit elliptical shapes throughout the observation period. The intuition behind this somewhat experimental approach is that the latent dimension d can be seen to measure the internal complexity of the observed multivariate time series, allowing us to identify from the smoothed curve intervals of time when the stocks behave in a more unified manner (low dimension) or independently of each other (high dimension). Note that the months close to the beginning and the end of the measurement interval belong on average to a fewer number of windows, meaning that they are expected to show more erratic behavior.

For simplicity, we apply the described procedure only with SURE2, using each of the four scatter matrices (and the two window lengths), and with Asymp, which has good overall performance in simulations when combined with the H–R estimator.

The results for SURE2 are shown in Fig. 6. In line with the evidence in Tsay (2010), we find at least two and in most time points at least three common features in our five stock case. The robust variants of SURE2 favour three (even four) dimensions, emphasizing also other common features than the market and industrial components. Interestingly, the time-varying patterns of the estimated dimensionality in the robust approaches seem to be largely in accordance with each other and following general market conditions where the major periods of decreasing prices (i.e. the beginning of 2000s and the financial crisis 2007–2009), and their onsets, are associated with somewhat lower dimensions. Finally, note that the fact that the curves for the non-robust “Cov” are markedly different from the robust ones is a clear indication that the data

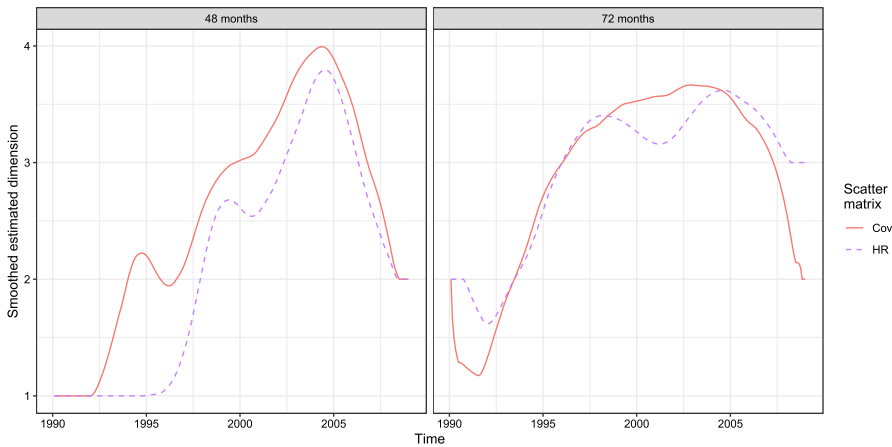


Fig. 7 The smoothed curves of latent dimensionalities estimated with Asymp using the window approach as described in the main text. The two panels correspond to the window lengths of $\ell = 48$ and $\ell = 72$ months, respectively

indeed exhibit heavy-tailed behaviour (as typically with asset returns) that hampers the estimation of the covariance matrix but leaves the robust estimates unaffected.

The results for Asymp(Cov) and Asymp(HR) are shown in Fig. 7 and match rather well with the corresponding results obtained with SURE2 in Fig. 6. For example, when using the window length of 72 months, both SURE2(HR) and Asymp(HR) identify 3–4 latent dimensions for the majority of the observation window, with two “bumps” in the curve. The only major difference is between SURE2(HR) and Asymp(HR) during the years 1990–2000 when the window length of 48 months is used: the former method claims a full set of 4 latent variables whereas the latter finds only a single one. This discrepancy is most likely explained by the short window length which is not large enough to estimate σ^2 (SURE2) or to invoke asymptotic arguments (Asymp). As such, it seems more reasonable to restrict attention to the window length of 72 months, where both methods agree that the amount of latent variables was lower at the start of the observation period and gradually increased with time.

7 Discussion

The results obtained in this work open up several avenues for future research. Perhaps most notably, our simulations revealed that the standard approach of choosing the parameter estimate to be the global minimizer of the SURE-criterion might not be optimal in the current scenario. As a simple alternative to minimization, we explored using a change point detection based approach which indeed proved superior to the minimization in various settings (and vice versa in other settings). As such, the matter clearly warrants more investigation. We also note that the dangers of “blindly” minimizing a model fit criterion are naturally well-known in the model selection literature.

Indeed, also Ulfarsson and Solo (2008) mention this caveat, although, not in the context of SURE but the Bayesian information criterion. Despite this, we are not aware of any general alternatives to minimization being proposed in the literature, discounting visual inspection (which can be seen as both heuristic and subjective). Quite possibly, such procedures may not even exist as their behaviour would depend greatly on the functional form of the particular criterion in question. But, at least in the current scenario of dimension estimation, the change point criterion appears to provide a feasible option.

A second point of interest brought up by our simulations is the sudden increase in the accuracy of the robust SURE2 in Fig. 4. In that particular data scenario there appears to be a “critical” sample size after which SURE2 achieves perfect estimation results. The dependence of this sample size on the model parameters, especially p and d is something to be studied and quantified. We note that any theoretical investigation of this matter is likely to be very difficult as it concerns the finite-sample properties of the method, whereas the large majority of the existing theoretical results in the robust literature are conducted in the asymptotic framework where $n \rightarrow \infty$.

As a third point, the simulations in Sect. 5.2 revealed that the Bayesian estimator by Minka (2000) gave a suprisingly good performance when combined with the H–R estimator of scatter. Such plug-in procedures in the context of Minka (2000) appear not to have been earlier investigated in the literature and clearly warrant future study. In the top panel of Fig. 3 the performance of the Bayesian plug-in estimator (Minka) is very similar to Asymp, meaning that for low n it appears to be an approximately equal to the asymptotical test, partially also explaining its bad performance for low n .

Fourthly, despite being a significant improvement over the Gaussian assumption, the elliptical model (1) can still be seen as somewhat limiting in practice. In particular, a consequence of the ellipticity is that the tails of the distribution are equally heavy along all directions in the space \mathbb{R}^p . One solution would be to consider the independent component model $x_i = \mu + Az_i$ instead, where $A \in \mathbb{R}^{p \times p}$ is full rank and the components of the p -vector z_i are mutually independent random variables (Comon and Jutten 2010). Exactly $p - d$ elements of z_i are assumed to be Gaussian and, as such, noise, making the signal dimension of the model equal to d . By taking the signal components of z_i to have different tail decays, a richer variety of heavy-tailed behaviours for x_i is obtained, see Virta et al. (2020). The independent component model admits a solution through the use of *pairs* of scatter functionals, see Tyler et al. (2009), and a similar approach could possibly serve as a starting point for deriving a SURE-based criterion for the estimation of the dimension d in the independent component model.

Fifthly, we comment on using the final eigenvalue s_p of S as the estimator of the noise variance. As mentioned earlier, we made this choice as it imposes minimal assumptions on the latent signal dimension d , allowing also cases where the ratio d/p is relatively large. This benefit can be seen in the top panel of Fig. 3, where robust SURE2 is able to estimate the latent dimension increasingly well even for very large values of d/p . The same would not be possible if, for example, the median or the third quartile of the eigenvalues would be used as the noise variance estimator. Similarly, in the asset return example in Sect. 6, the ratio d/p was estimated to be larger than 0.6 in majority of the time windows. And although our example data were rather

low-dimensional, large values of d/p are known to occur also in higher-dimensional economic data sets, making the use of estimators such as s_p recommended in such contexts.

On the flipside, using the estimator s_p most likely compromises the unbiasedness of our risk estimate. That is, even though $E(R_{2k}) = (1/n) \sum_{i=1}^n E\|\hat{x}_i - V_0 y_i\|^2$, there is no guarantee that the plug-in risk \hat{R}_{2k} would be an unbiased estimator of the risk. We investigated this with some preliminary simulations (not shown here) and were led to two conclusions: (i) It appears that \hat{R}_{2k} is indeed biased but that the bias vanishes when $n \rightarrow \infty$. Thus, the plugging in of s_p likely changes our risk estimator from unbiased to asymptotically unbiased. (ii) Using the “oracle” estimator R_{2k} (where the true value of σ^2 is used) in place of \hat{R}_{2k} did not yield significantly better results and, in fact, in some cases \hat{R}_{2k} performed strictly better than R_{2k} . The effect of estimating σ^2 clearly warrants further study but, based on the previous, it does not appear that its estimation affects the performance of the method too much.

As pointed out to us by an anonymous reviewer, an alternative approach to the dimension estimation problem would be to instead minimize the risk function

$$E\left(\|t_0 + P_k(\tilde{x} - t_0) - V_0 \tilde{y}\|^2\right), \quad (9)$$

where the location and projection estimates t_0 , P_k are estimated based on the sample x_1, \dots, x_n and \tilde{x} (and the corresponding latent signal \tilde{y}) is another draw from the model, independent of the actual sample. Thus, the difference to our proposed concept is that in (9) an independent “test sample” \tilde{x} is used to evaluate the error in the PCA-reconstruction. While this is a perfectly valid approach to the problem, it no longer adheres to the SURE-framework (but is more akin to the classical Akaike’s information criterion) and we leave its study to future work.

Finally, an interesting alternative to our model (1) would be the elliptical factor model $x_i = \mu + V_0 D y_i + \varepsilon_i$ where $V_0 \in \mathbb{R}^{p \times d}$ has orthonormal columns, D is diagonal and $y_i \in \mathbb{R}^d$ and $\varepsilon_i \in \mathbb{R}^p$ are mutually independent and spherical. As general scatter matrices do not have the additivity property possessed by the covariance matrix, it is not guaranteed that the methodology used in this work would be guaranteed to identify the latent dimension of this model via the “eigengap”. Our preliminary investigations however give promising results regarding this and reveal that the estimation can indeed work, but due to the lack of theoretical motivation, we have left this too for future work.

Acknowledgements The work of Niko Lietzén and Henri Nyberg was supported by the Academy of Finland (Grant 321968). The work of Joni Virta was supported by the Academy of Finland (Grants 335077, 347501 and 353769). Nyberg also acknowledges financial support by the Emil Aaltonen Foundation (Grant 180287) and the Foundation for Economic Education (Liikesivistysrahasto, Grant 220246).

Funding Open Access funding provided by University of Turku (including Turku University Central Hospital).

Declarations

Conflict of interest The authors state that they have no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Proofs of the technical results

Proof of Lemma 1 We write,

$$\begin{aligned} \mathbb{E}\|\hat{x}_i - V_0 y_i\|^2 &= \mathbb{E}\|(\hat{x}_i - x_i) + (x_i - V_0 y_i)\|^2 \\ &= \mathbb{E}\|\hat{x}_i - x_i\|^2 + 2\mathbb{E}\{(\hat{x}_i - x_i)' \varepsilon_i\} + \mathbb{E}\|\varepsilon_i\|^2. \end{aligned}$$

Now, $\|\hat{x}_i - x_i\|^2 = \text{tr}\{(I_p - P_k)(x_i - t_0)(x_i - t_0)'\}$, implying that averaging over i in the preceding expansion yields

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}\|\hat{x}_i - V_0 y_i\|^2 = \mathbb{E}[\text{tr}\{(I_p - P_k)S_0\}] + \frac{2}{n} \sum_{i=1}^n \mathbb{E}(\hat{x}'_i \varepsilon_i) - p\sigma^2.$$

Hence, the claim follows as soon as we show that $\mathbb{E}(\hat{x}_{ij} \varepsilon_{ij}) = \sigma^2 \mathbb{E}\{(\partial/\partial x_{ij})\hat{x}_{ij}\}$. To see this, we use the law of total expectation in conjunction with Stein’s lemma to write,

$$\begin{aligned} \mathbb{E}(\hat{x}_{ij} \varepsilon_{ij}) &= \mathbb{E}\{\mathbb{E}(\hat{x}_{ij} \varepsilon_{ij} \mid y_1, \dots, y_n)\} \\ &= \sigma^2 \mathbb{E}[\mathbb{E}\{(\partial/\partial \varepsilon_{ij})\hat{x}_{ij} \mid y_1, \dots, y_n\}] \\ &= \sigma^2 \mathbb{E}[\mathbb{E}\{(\partial/\partial x_{ij})\hat{x}_{ij} \mid y_1, \dots, y_n\}] \\ &= \sigma^2 \mathbb{E}\{(\partial/\partial x_{ij})\hat{x}_{ij}\}, \end{aligned}$$

where the second equality uses Stein’s lemma on the reconstruction \hat{x}_{ij} treated as a function of the full data (which are, conditional on y_1, \dots, y_n , Gaussian), the third equality uses the fact that x_{ij} and ε_{ij} are equal up to translation by a constant (again, under the conditioning).

Proof of Lemma 3 We use, for brevity, the notation $t := t(F_n)$, $t^* := t(F_{n,i,j,\varepsilon})$, and similarly for other quantities. Hence, by Assumption 1, we have $t^* = t + \varepsilon h_{ij} + o(\varepsilon)$ and $S^* = S + \varepsilon H_{ij} + o(\varepsilon)$. Thus, by [Stewart (2001) Chapter 1.3.2], the matrix S^* has a ℓ th eigenvector u_ℓ^* with the following first-order expansion,

$$u_\ell^* = u_\ell + \varepsilon \sum_{m \neq \ell} \frac{1}{s_\ell - s_m} T_m H_{ij} u_\ell + o(\varepsilon), \tag{A1}$$

where u_ℓ is the ℓ th eigenvector of S and $T_\ell = u_\ell u_\ell'$ is the corresponding orthogonal projection. Taking outer products on both side of (A1) gives

$$T_\ell^* = T_\ell + \varepsilon \sum_{m \neq \ell}^p \frac{1}{s_\ell - s_m} (T_\ell H_{ij} T_m + T_m H_{ij} T_\ell) + o(\varepsilon).$$

Consequently, the projection P_k^* onto the space spanned by the eigenvectors corresponding to the k largest eigenvalues of S^* (which are, for small enough ε , simple) is $P_k^* = P_k + \varepsilon A_{ij} + o(\varepsilon)$, where A_{ij} is as in the statement of the lemma (note that, by symmetry of the H_{ij} , all terms with $m \leq k$ get cancelled).

Now, recall that $x_i^* = x_i + \varepsilon e_j$ and write,

$$\begin{aligned} \hat{x}_i^* - \hat{x}_i &= t^* - t + P_k^*(x_i^* - t^*) - P_k(x_i - t) \\ &= \varepsilon \{h_{ij} + P_k(e_j - h_{ij}) + A_{ij}(x_i - t)\} + o(\varepsilon), \end{aligned}$$

showing that the desired derivative equals

$$\frac{\partial}{\partial x_{ij}} \hat{x}_{ij} = e_j' h_{ij} + e_j' P_k(e_j - h_{ij}) + e_j' A_{ij}(x_i - t).$$

Summation over j now yields the claim. □

Proof of Lemma 4 We write $t := t(F_n)$ and $t^* := t(F_{n,i,j,\varepsilon})$, and similarly for S and S^* . Clearly, $t^* = t + (1/n)\varepsilon e_j$, implying that $x_i + \varepsilon e_j - t^* = x_i - t + (1 - 1/n)\varepsilon e_j$. Writing $y_i := x_i - t$, we get,

$$\begin{aligned} S^* &= \frac{1}{n} \sum_{\ell \neq i}^n \{y_\ell - (1/n)\varepsilon e_j\} \{y_\ell - (1/n)\varepsilon e_j\}' \\ &\quad + \frac{1}{n} \{y_i + (1 - 1/n)\varepsilon e_j\} \{y_i + (1 - 1/n)\varepsilon e_j\}'. \end{aligned} \tag{A2}$$

The first term of (A2) simplifies to

$$S - \frac{1}{n} y_i y_i' + \frac{1}{n^2} \varepsilon e_j y_i' + \frac{1}{n^2} \varepsilon y_i e_j' + o(\varepsilon),$$

whereas the second term can be written as

$$\frac{1}{n} y_i y_i' + \frac{n-1}{n^2} \varepsilon e_j y_i' + \frac{n-1}{n^2} \varepsilon y_i e_j' + o(\varepsilon).$$

Combining these two now yields the claim. □

Proof of Lemma 5 We begin with h_{ij} . Let $g : \mathbb{R}^p \rightarrow \mathbb{R}$ be the objective function with $g(t) = (1/n) \sum_{i=1}^n \|x_i - t\|$, from whose minimization the spatial median $t(F_n)$ is found. It is straightforwardly checked that g is convex and, since, by the assumption

that $t(F_n) \neq x_i$ for all $i = 1, \dots, n$, g is differentiable in a neighbourhood of $t(F_n)$, the gradient of g must vanish at $t(F_n)$,

$$\frac{1}{n} \sum_{i=1}^n \frac{x_i - t(F_n)}{\|x_i - t(F_n)\|} = 0. \tag{A3}$$

Define now the function $f : \mathbb{R}^{p+1} \rightarrow \mathbb{R}^p$ such that

$$f(\varepsilon, t) = \frac{1}{n} \sum_{\ell \neq i} \frac{x_i - t}{\|x_i - t\|} + \frac{1}{n} \frac{x_i + \varepsilon e_j - t}{\|x_i + \varepsilon e_j - t\|}.$$

Now, $f\{0, t(F_n)\} = 0$ and, moreover, f is continuously differentiable at $(0, t(F_n))$ and has the Jacobian $D_t f\{0, t(F_n)\} = -(1/n)G$ where G is as defined in the statement of the lemma. Now, G is a sum of weighted projection matrices to the orthogonal complements of the lines spanned by the vectors $y_i = x_i - t(F_n)$. Hence, G has full rank (by our assumption that the points x_1, \dots, x_n are not concentrated on a line) and the implicit function theorem then guarantees that for a suitably small neighbourhood $\mathcal{S} \subset \mathbb{R}$ of zero there exists a differentiable function $b : \mathcal{S} \rightarrow \mathbb{R}^p$ such that $f\{\varepsilon, b(\varepsilon)\} = 0$.

In the following, with some abuse of notation, \mathcal{S} will denote a (changing) small enough neighbourhood of zero in \mathbb{R} . By our sample being not concentrated on a line and Milasevic and Ducharme (1987), the spatial median of $F_{n,i,j,\varepsilon}$ is unique in \mathcal{S} . Let $M := \max_i \|x_i\|$ and take any $t \in \mathbb{R}^p$ such that $\|t\| \geq 3M + 1$. Then, $(1/n) \sum_{i=1}^n \|x_i - t\| \geq 2M + 1$, whereas $(1/n) \sum_{i=1}^n \|x_i - 0\| \leq M$, showing that, for all $\varepsilon \in \mathcal{S}$, the spatial medians $t(F_{n,i,j,\varepsilon})$ reside in a compact set. Thus, by Berge’s Maximum Theorem, the map $\varepsilon \mapsto t(F_{n,i,j,\varepsilon})$ is continuous in \mathcal{S} , implying that the equivalent of the assumption that $t(F_n) \neq x_i$, for all $i = 1, \dots, n$, holds for $F_{n,i,j,\varepsilon}$. Consequently, for any $\varepsilon \in \mathcal{S}$, the corresponding spatial median objective function is differentiable and its gradient vanishes at the spatial median $t(F_{n,i,j,\varepsilon})$. Its gradient is equal to $f(\varepsilon, t)$ and, by the uniqueness of the spatial medians $t(F_{n,i,j,\varepsilon})$, we thus have that the implicit function b actually traces the spatial medians as a function of ε , i.e., $t(F_{n,i,j,\varepsilon}) = b(\varepsilon)$ for $\varepsilon \in \mathcal{S}$.

The differentiability of b at zero now entails that we may write $t(F_{n,i,j,\varepsilon}) = b(\varepsilon) = t(F_n) + \varepsilon h_{ij} + o(\varepsilon)$ for some $h_{ij} \in \mathbb{R}^p$ as $\varepsilon \rightarrow 0$. Plugging the expansion in to the $F_{n,i,j,\varepsilon}$ -equivalent of (A3), we get

$$0 = \sum_{\ell \neq i} \frac{y_\ell - \varepsilon h_{ij} + o(\varepsilon)}{\|y_\ell - \varepsilon h_{ij} + o(\varepsilon)\|} + \frac{y_i + \varepsilon(e_j - h_{ij}) + o(\varepsilon)}{\|y_i + \varepsilon(e_j - h_{ij}) + o(\varepsilon)\|}, \tag{A4}$$

where $y_i = x_i - t(F_n)$. The first term in (A4) splits into two parts, the first of which simplifies as,

$$\sum_{\ell \neq i} \frac{y_\ell}{\|y_\ell - \varepsilon h_{ij} + o(\varepsilon)\|} = \sum_{\ell \neq i} \frac{y_\ell}{\|y_\ell\|} + \varepsilon \sum_{\ell \neq i} \frac{y_\ell y'_\ell}{\|y_\ell\|^3} h_{ij} + o(\varepsilon),$$

whereas, the second part writes,

$$-\sum_{\ell \neq i}^n \frac{\varepsilon h_{ij} + o(\varepsilon)}{\|y_\ell - \varepsilon h_{ij} + o(\varepsilon)\|} = -\varepsilon \sum_{\ell \neq i}^n \frac{1}{\|y_\ell\|} h_{ij} + o(\varepsilon).$$

Simplifying the second term of (A4) similarly yields,

$$\frac{y_i + \varepsilon(e_j - h_{ij}) + o(\varepsilon)}{\|y_i + \varepsilon(e_j - h_{ij}) + o(\varepsilon)\|} = \frac{y_i}{\|y_i\|} - \varepsilon \frac{y_i y_i'}{\|y_i\|^3} (e_j - h_{ij}) + \varepsilon \frac{1}{\|y_i\|} (e_j - h_{ij}) + o(\varepsilon).$$

Plugging everything now back in to (A4), and using (A3), we get, in the notation of the statement of the lemma,

$$-\varepsilon G h_{ij} + \varepsilon A_i e_j + o(\varepsilon) = 0.$$

Division by ε and letting $\varepsilon \rightarrow 0$ then yields the desired expression for h_{ij} .

For H_{ij} , resorting to the same notation as in the proof of Lemma 3 and writing $y_i^* := x_i^* - t^*$, we first have,

$$\|y_i^*\|^2 - \|y_i\|^2 = 2\varepsilon(e_j - h_{ij})' y_i + o(\varepsilon),$$

and,

$$y_i^* (y_i^*)' - y_i y_i' = \varepsilon \{ (e_j - h_{ij}) y_i' + y_i (e_j - h_{ij})' \} + o(\varepsilon) =: \varepsilon B_{ij} + o(\varepsilon).$$

For $\ell \neq i$, we also write $y_\ell^* := x_\ell - t^*$ and have

$$\|y_\ell^*\|^2 - \|y_\ell\|^2 = -2\varepsilon h_{ij}' y_\ell + o(\varepsilon),$$

and

$$y_\ell^* (y_\ell^*)' - y_\ell y_\ell' = -\varepsilon \{ h_{ij} y_\ell' + y_\ell h_{ij}' \} + o(\varepsilon) =: -\varepsilon C_{\ell ij} + o(\varepsilon).$$

Hence, the partition

$$S^* - S = \frac{1}{n} \sum_{\ell \neq i}^n \left\{ \frac{y_\ell^* (y_\ell^*)'}{\|y_\ell^*\|^2} - \frac{y_\ell y_\ell'}{\|y_\ell\|^2} \right\} + \frac{1}{n} \left\{ \frac{y_i^* (y_i^*)'}{\|y_i^*\|^2} - \frac{y_i y_i'}{\|y_i\|^2} \right\},$$

along with the formula for general matrices A, B and scalars a, b with $a \neq 0$ that,

$$\frac{A + \varepsilon B + o(\varepsilon)}{a + \varepsilon b + o(\varepsilon)} = \frac{A}{a} + \varepsilon \frac{aB - bA}{a^2} + o(\varepsilon),$$

implies that

$$S^* - S = -\varepsilon \frac{1}{n} \sum_{\ell=1}^n \{ \|y_\ell\|^{-2} (h_{ij} y'_\ell + y_\ell h'_{ij}) - 2 \|y_\ell\|^{-4} y_\ell h'_{ij} y_\ell y'_\ell \} \\ + \varepsilon \frac{1}{n} \{ \|y_i\|^{-2} (e_j y'_i + y_i e'_j) - 2 \|y_i\|^{-4} y_i e'_j y_i y'_i \} + o(\varepsilon).$$

Simplifying the expression now yields the claim. \square

References

- Anderson TW (1963) Asymptotic theory for principal component analysis. *Ann Math Stat* 34(1):122–148
- Bernard G, Verdebout T (2021) On some multivariate sign tests for scatter matrix eigenvalues. *Economet Stat*
- Borak S, Misiorek A, Weron R (2011) Models for heavy-tailed asset returns. In: *Statistical tools for finance and insurance*, pp 21–55. Springer, Berlin
- Brown B (1983) Statistical uses of the spatial median. *J R Stat Soc Ser B* 45(1):25–30
- Comon P, Jutten C (2010) *Handbook of blind source separation: independent component analysis and applications*. Academic Press, Cambridge
- Croux C, Haesbroeck G (2000) Principal component analysis based on robust estimators of the covariance or correlation matrix: influence functions and efficiencies. *Biometrika* 87(3):603–618
- Deng WQ, Craiu RV (2023) Exploring dimension learning via a penalized probabilistic principal component analysis. *J Stat Comput Simul* 93(2):266–297
- Dümbgen L, Tyler DE (2005) On the breakdown properties of some multivariate M-functionals. *Scand J Stat* 32(2):247–264
- Dürre A, Tyler DE, Vogel D (2016) On the eigenvalues of the spatial sign covariance matrix in more than two dimensions. *Stat Prob Lett* 111:80–85
- Fan J, Wang W, Zhu Z (2021) A shrinkage principle for heavy-tailed data: high-dimensional robust low-rank matrix recovery. *Ann Stat* 49(3):1239
- Fang KW (2018) *Symmetric multivariate and related distributions*. CRC Press, New York
- Gai J, Stevenson RL (2010) Studentized dynamical system for robust object tracking. *IEEE Trans Image Process* 20(1):186–199
- Gai J, Li Y, Stevenson RL (2008) An EM algorithm for robust Bayesian PCA with Student's t-distribution. In: *2008 15th IEEE International Conference on Image Processing*, pp. 2672–2675 . IEEE
- Haldane J (1948) Note on the median of a multivariate distribution. *Biometrika* 35(3–4):414–417
- Hettmansperger TP, Randles RH (2002) A practical affine equivariant multivariate median. *Biometrika* 89(4):851–860
- Killick R, Eckley I (2014) changepoint: an R package for changepoint analysis. *J Stat Softw* 58(3):1–19
- Luo W, Li B (2016) Combining eigenvalues and variation of eigenvectors for order determination. *Biometrika* 103(4):875–887
- Luo W, Li B (2021) On order determination by predictor augmentation. *Biometrika* 108(3):557–574
- Magyar A, Tyler DE (2011) The asymptotic efficiency of the spatial median for elliptically symmetric distributions. *Sankhya B* 73(2):165–192
- Marden JI (1999) Some robust estimates of principal components. *Stat Prob Lett* 43(4):349–359
- Milasevic P, Ducharme G (1987) Uniqueness of the spatial median. *Ann Stat* 15(3):1332–1333
- Minka T (2000) Automatic choice of dimensionality for PCA. *Adv Neural Inform Process Syst* 78:13
- Nordhausen K, Tyler DE (2015) A cautionary note on robust covariance plug-in methods. *Biometrika* 102(3):573–588
- Nordhausen K, Oja H, Tyler DE (2021) Asymptotic and bootstrap tests for subspace dimension. *J Multivar Anal* 8:104830 104830
- Nordhausen K, Oja H, Tyler DE, Virta J (2021) ICtest: estimating and testing the number of interesting components in linear dimension reduction. R package version 0.3-4

- Nordhausen K, Sirkia S, Oja H, Tyler DE (2018) ICSNP: Tools for multivariate nonparametrics. R package version 1.1-1. <https://CRAN.R-project.org/package=ICSNP>
- Oja H (2010) Multivariate nonparametric methods with R: an approach based on spatial signs and ranks. Springer, New York
- Paindaveine D (2008) A canonical definition of shape. *Stat Prob Lett* 78(14):2240–2247
- Pison G, Rousseeuw PJ, Filzmoser P, Croux C (2003) Robust factor analysis. *J Multivar Anal* 84(1):145–172
- Schott JR (2006) A high-dimensional test for the equality of the smallest eigenvalues of a covariance matrix. *J Multivar Anal* 97(4):827–843
- Stein CM (1981) Estimation of the mean of a multivariate normal distribution. *Ann Stat* 8:1135–1151
- Stewart GW (2001) Matrix algorithms: volume ii: eigensystems. SIAM, Philadelphia
- Tibshirani RJ, Taylor J (2012) Degrees of freedom in lasso problems. *Ann Stat* 40(2):1198–1232
- Tipping ME, Bishop CM (1999) Probabilistic principal component analysis. *J R Stat Soc Ser B* 61(3):611–622
- Tsay RS (2010) Analysis of financial time series. Wiley, Hoboken
- Tyler DE (1987) A distribution-free M-estimator of multivariate scatter. *Ann Stat* 7:234–251
- Tyler DE, Critchley F, Dümbgen L, Oja H (2009) Invariant co-ordinate selection. *J R Stat Soc Ser B* 71(3):549–592
- Ulfarsson MO, Solo V (2008) Dimension estimation in noisy PCA with SURE and random matrix theory. *IEEE Trans Signal Process* 56(12):5804–5816
- Ulfarsson MO, Solo V (2015) Selecting the number of principal components with SURE. *IEEE Signal Process Lett* 22(2):239–243
- Virta J, Lietzén N, Viitasari L, Ilmonen P (2020) Latent model extreme value index estimation. arXiv preprint [arXiv:2003.10330](https://arxiv.org/abs/2003.10330)
- Visuri S, Koivunen V, Oja H (2000) Sign and rank covariance matrices. *J Stat Plan Inference* 91(2):557–575
- Vogel D, Fried R (2015) Robust change detection in the dependence structure of multivariate time series. *Modern Nonparametric. robust and multivariate methods*. Springer, Cham, pp 265–288
- Wax M, Kailath T (1985) Detection of signals by information theoretic criteria. *IEEE Trans Acoust Speech Signal Proces* 33(2):387–392
- Wiesel A (2012) Geodesic convexity and covariance estimation. *IEEE Trans Signal Process* 60(12):6182–6189
- Zhao L, Krishnaiah PR, Bai Z (1986) On detection of the number of signals in presence of white noise. *J Multivar Anal* 20(1):1–25

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.