

SELECTING THE NUMBER OF PRINCIPAL COMPONENTS: ESTIMATION OF THE TRUE RANK OF A NOISY MATRIX

BY YUNJIN CHOI*, JONATHAN TAYLOR^{†,1} AND ROBERT TIBSHIRANI^{†,2}

National University of Singapore and Stanford University[†]*

Principal component analysis (PCA) is a well-known tool in multivariate statistics. One significant challenge in using PCA is the choice of the number of principal components. In order to address this challenge, we propose distribution-based methods with *exact* type 1 error controls for hypothesis testing and construction of confidence intervals for signals in a noisy matrix with finite samples. Assuming Gaussian noise, we derive exact type 1 error controls based on the conditional distribution of the singular values of a Gaussian matrix by utilizing a post-selection inference framework, and extending the approach of [Taylor, Loftus and Tibshirani (2013)] in a PCA setting. In simulation studies, we find that our proposed methods compare well to existing approaches.

1. Introduction.

1.1. *Overview.* Principal component analysis (PCA) is a commonly used method in multivariate statistics. It can be used for a variety of purposes including as a descriptive tool for examining the structure of a data matrix, as a pre-processing step for reducing the dimension of the column space of the matrix [Josse and Husson (2012)], or for matrix completion [Cai, Candès and Shen (2010)].

One important challenge in PCA is how to determine the number of principal components to retain. Jolliffe (2002) provides an excellent summary of existing approaches to determining the number of principal components, grouping them into three branches: subjective methods (e.g., the scree plot), distribution-based test tools (e.g., Bartlett’s test) and computational procedures (e.g., cross-validation). Each branch has advantages as well as disadvantages, and no single method has emerged as the community standard.

Figure 1 offers a scree plot as an example. The data are five test scores from 88 students [taken from Mardia, Kent and Bibby (1979), pages 3–4]; the figure shows the five singular values in decreasing order. The “elbow” in this plot seems to occur at rank two or three, but it is not clearly visible. We revisit this example with our proposed approach in Section 5.3.

Received May 2015; revised November 2016.

¹Supported by NSF Grant DMS-12-08857 and AFOSR Grant 113039.

²Supported by NSF Grant DMS-99-71405 and NIH Contract N01-HV-28183.
MSC2010 subject classifications. 62J05, 62J07, 62F03.

Key words and phrases. Principal components, hypothesis test, exact p -value.

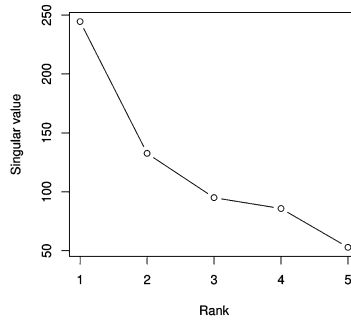


FIG. 1. Singular values of the score data in decreasing order. The data consist of exam scores of 88 students on five different topics (Mechanics, Vectors, Algebra, Analysis and Statistics).

In this paper, we propose a class of statistical methods utilizing hypothesis testing framework for determining the rank of the signal matrix in a noisy matrix model. The estimated rank here corresponds to the number of principal components to retain in PCA. Under Gaussian assumption, the proposed hypothesis testing method provides exact type 1 error controls along with exact confidence intervals of signal parameters in finite samples.

1.2. *Related works.* Our approach is analogous to the exact test of variables in forward stepwise regression proposed by Tibshirani et al. (2014). Both approaches sequentially investigate statistical significance of the statistic in interest—principal components or variables, and provide exact type 1 error controls via conditioning on the selection event. To achieve exact type 1 error control in forward stepwise regression, Tibshirani et al. (2014) proposed the truncated Gaussian test conditioning on the event of selecting active variables. The proposed truncated Gaussian test utilizes the fact that the selection event can be characterized by an observed data vector y falling into a polyhedral set. Despite the analogous objectives of PCA and regression, the truncated Gaussian test is not applicable to PCA; in PCA, the event of selecting principal components cannot be characterized by an observed data matrix Y falling into a polyhedral set. While selecting a variable is a discrete event in forward stepwise regression, a principal components in PCA are chosen from a continuum for Y being a matrix. As the domain of the selection event is a continuum, the resulting null distribution conditional on the the selection event is defined on a measure zero domain rather than being a truncated Gaussian distribution. Therefore, establishing an exact test in PCA via conditioning requires subtle approach.

To address this challenge, we utilize the *Kac–Rice test*, an exact method for testing and constructing confidence intervals for signals under a global null hypothesis in adaptive regression. Under the global null scenario, one of our proposed methods corresponds to the application of the *Kac–Rice test* to a penalized regression minimizing the Frobenius norm with a nuclear norm penalty. In this paper, we

extend the *Kac–Rice test* and the construction of confidence intervals to not only the global null scenario but the general case. Also in the global null scenario, one of the extended methods provides stronger power than the *Kac–Rice test*. The exact property of the *Kac–Rice test* is preserved in extension to a general step by incorporating a post-selection inference framework. The resulting statistics use a conditional survival function of the eigenvalues of a Wishart matrix.

In the context of inference based on the distribution of eigenvalues, [Muirhead \(1982\)](#), Theorem 9.6.2 on page 409, and, more recently, [Kritchman and Nadler \(2008\)](#) have proposed methods for testing essentially the same hypothesis as in this paper. Both [Muirhead \(1982\)](#), Theorem 9.6.2 on page 409, and [Kritchman and Nadler \(2008\)](#) benefit from an asymptotic distribution of the test statistic: [Muirhead \(1982\)](#), Theorem 9.6.2 on page 409, forms a likelihood ratio test with the asymptotic Chi-square distribution, and [Kritchman and Nadler \(2008\)](#) use the Tracy–Widom law, which is the asymptotic distribution of the largest eigenvalue of a Wishart matrix, incorporating the result of [Johnstone \(2001\)](#). While these test methods show conservative results and thus lose signal detection power in the general stage, our proposed methods provide exact type 1 error controls and decent detection power at the same time. Additionally, our approach provides a method for constructing confidence intervals in addition to hypothesis testing.

1.3. *Organization of the paper.* The rest of the paper is organized as follows. In Section 2, we propose methods based on the distribution of eigenvalues of a Wishart matrix. Section 2.1 introduces the main points of the *Kac–Rice test* [[Taylor, Loftus and Tibshirani \(2013\)](#)] from which we derive our proposals. Then we propose hypothesis testing procedures regarding the true rank of a signal matrix in Section 2.2. Our method for constructing exact confidence intervals of signals is described in Section 2.3.

In Section 3, we propose a method for estimating the rank of a signal matrix from the hypothesis tests. We illustrate a sequential hypothesis testing procedure for determining the matrix rank, based on a proposal of [G’Sell et al. \(2013\)](#).

Section 4 introduces a data-driven method for estimating the noise level. Our proposed method is analogous to cross-validation. In order to apply cross-validation to a data matrix, we use a method proposed by [Mazumder, Hastie and Tibshirani \(2010\)](#).

Section 5 provides additional examples of the proposed methods. In Section 5.1, simulation results with estimated noise level are presented. Section 5.2 shows simulation results with non-Gaussian noise to check for robustness. We revisit the real data example introduced in Figure 1 in Section 5.3. The paper concludes with a brief discussion in Section 6.

2. Proposed distribution-based methods. Throughout this paper, we assume that the observed data matrix $Y \in \mathbb{R}^{N \times p}$ is the sum of a low-rank signal

matrix $B \in \mathbb{R}^{N \times p}$ and a Gaussian noise matrix $E \in \mathbb{R}^{N \times p}$ as follows:

$$\begin{aligned}
 Y &= B + E, \\
 \text{rank}(B) &= \kappa < \min(N, p), \\
 E_{ij} &\sim \text{i.i.d. } N(0, \sigma^2) \quad \text{for } i \in \{1, \dots, N\}, j \in \{1, \dots, p\}
 \end{aligned}$$

that is,

$$(2.1) \quad Y \sim N(B, \sigma^2 I_N \otimes I_p).$$

Here, following the notation from Muirhead (1982), page 73, $Y \sim N(B, \sigma^2 I_N \otimes I_p)$ in (2.1) denotes

$$\begin{pmatrix} Y_1 \\ \vdots \\ Y_N \end{pmatrix} \sim N \left(\begin{pmatrix} B_1 \\ \vdots \\ B_N \end{pmatrix}, \begin{pmatrix} \sigma^2 I_p & 0 & \cdots & 0 \\ 0 & \sigma^2 I_p & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma^2 I_p \end{pmatrix} \right),$$

where Y_j and B_j represent the j th row vectors of Y and B , respectively, and the Kronecker product is denoted by \otimes . In this paper, without loss of generality, we assume that $N > p$. For the case of $N < p$, we can simply switch the role of N and p by dealing with Y^T instead of Y . In this paper, we focus on finding κ , the rank of the signal matrix B , and the construction of confidence intervals for the signals in B .

With a centered data matrix, our proposed approaches with this model assumption are valid for the popular spiked covariance model as well; both traditional spiked covariance approaches and our proposed tests with a centered data matrix are examining basically the same statistics, noting that the centering operation of a data matrix does not change the rank of B . The statistics in interest for the traditional spiked covariance model approaches are the eigenvalues of a covariance matrix, and our proposed methods investigate the singular values of a data matrix, where the eigenvalues from the covariance matrix are the same as the squared singular values of a centered data matrix.

We first review the global null test and confidence interval construction of the first signal of Taylor, Loftus and Tibshirani (2013) and its application in matrix denoising problem in Section 2.1 which corresponds to the PCA setting. Then we extend the global null test to a general test procedure for testing $H_{k,0} : \text{rank}(B) \leq k - 1$ versus $H_{k,1} : \text{rank}(B) \geq k$ for $k = 1, \dots, p - 1$ in Section 2.2 and describe how to construct confidence intervals for the k th largest signal parameters in Section 2.3.

2.1. *Review of the Kac–Rice test.* We briefly discuss the framework of Taylor, Loftus and Tibshirani (2013) and its application to a matrix denoising problem,

which we extend further later in this paper. Section 2.1.1 covers the testing procedure for a class of null hypothesis. This null hypothesis corresponds to the global null in our matrix denoising problem. In Section 2.1.2, we discuss the construction of a confidence interval for the largest signal.

2.1.1. *Global null hypothesis testing.* Taylor, Loftus and Tibshirani (2013) derived the *Kac–Rice test*, a test providing exact type 1 error controls for a class of regularized regression problems of the following form:

$$(2.2) \quad \hat{\beta} \in \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \cdot \mathcal{P}(\beta)$$

with an outcome $y \in \mathbb{R}^p$, a predictor matrix $X \in \mathbb{R}^{N \times p}$ and a penalty term $\mathcal{P}(\cdot)$ with a regularization parameter $\lambda \geq 0$. Assuming that the outcome $y \in \mathbb{R}^N$ is generated from

$$y \sim N(X\beta, \Sigma),$$

the *Kac–Rice test* [Taylor, Loftus and Tibshirani (2013)] provides a method for testing

$$(2.3) \quad H_0 : \mathcal{P}(\beta) = 0$$

that yields exact type 1 error controls under the assumption that the penalty function \mathcal{P} is a support function of a convex set $\mathcal{C} \subseteq \mathbb{R}^p$, that is,

$$\mathcal{P}(\beta) = \max_{u \in \mathcal{C}} u^T \beta.$$

When applied to a matrix denoising problem of a popular form, (2.3) becomes a global null hypothesis:

$$(2.4) \quad H_0 : \Lambda_1 = 0 \equiv \operatorname{rank}(B) = 0 \equiv B = 0_{N \times p},$$

where $\Lambda_1 \geq \Lambda_2 \geq \dots \geq \Lambda_p \geq 0$ denote the singular values of B . Here are the details. For an observed data matrix $Y \in \mathbb{R}^{N \times p}$, a widely used method to recover the signal matrix B in (2.1) is to solve the following criterion:

$$(2.5) \quad \hat{B} \in \operatorname{argmin}_{B \in \mathbb{R}^{N \times p}} \frac{1}{2} \|Y - B\|_F^2 + \lambda \|B\|_* \quad \text{where } \lambda > 0,$$

where $\|\cdot\|_F$ and $\|\cdot\|_*$ denote a Frobenius norm and a nuclear norm, respectively. The nuclear norm term plays an analogous role as an ℓ_1 penalty term in lasso regression [Tibshirani (1996)]. The objective function (2.5) falls into the class of regression problems described in (2.2), with the predictor matrix X being $I_N \otimes I_p$ and the penalty function $\mathcal{P}(\cdot)$ being

$$\mathcal{P}(B) = \|B\|_* = \max_{u \in \mathcal{C}} \langle u, B \rangle$$

with $\mathcal{C} = \{A : \|A\|_{op} \leq 1\}$ where $\|\cdot\|_{op}$ denotes a spectral norm. We can therefore directly apply the *Kac–Rice test* with the resulting test statistic as follows, under the assumed model discussed in the beginning of Section 2:

$$(2.6) \quad \mathbb{S}_{1,0} = \frac{\int_{d_1}^{\infty} e^{-\frac{z^2}{2\sigma^2}} z^{N-p} \prod_{j=2}^p (z^2 - d_j^2) dz}{\int_{d_2}^{\infty} e^{-\frac{z^2}{2\sigma^2}} z^{N-p} \prod_{j=2}^p (z^2 - d_j^2) dz},$$

where $d_1 \geq d_2 \geq \dots \geq d_p \geq 0$ denote the observed singular values of Y . The test statistic $\mathbb{S}_{1,0}$ in (2.6) is uniformly distributed under the null hypothesis (2.4) and provides exact type 1 error controls for testing the global null hypothesis. Here, the value $\mathbb{S}_{1,0}$ represents the probability of observing more extreme values than d_1 under the null hypothesis, which coincides with the traditional notion of p -value.

Viewed differently, the test statistic $\mathbb{S}_{1,0}$ corresponds to a conditional survival probability of the largest observed singular value d_1 conditioned on all the other observed singular values d_2, \dots, d_p . The integrand of $\mathbb{S}_{1,0}$ coincides with the conditional distribution of the largest eigenvalue of a central Wishart matrix [James (1964)] via a change of variables. The denominator of $\mathbb{S}_{1,0}$ acts as a normalizing constant since the domain of the largest singular value d_1 conditioning on all the other singular values becomes (d_2, ∞) . A small magnitude of $\mathbb{S}_{1,0}$ implies large d_1 compared to d_2 , and thus supports $H_1 : \Lambda_1 > 0$.

2.1.2. *Confidence intervals for the largest signal.* Along with the *Kac–Rice test* mentioned in Section 2.1.1, a procedure for constructing an exact confidence interval for the leading signal in adaptive regression is proposed in Taylor, Loftus and Tibshirani (2013). As in Section 2.1.1, by applying the result of Taylor, Loftus and Tibshirani (2013) to our matrix denoising setting, we can generate an exact confidence interval for $\tilde{\Lambda}_1$ which is defined as follows:

$$\tilde{\Lambda}_1 = \langle U_1 V_1^T, B \rangle,$$

where $Y = UDV^T$ is a singular value decomposition of Y with $D = \text{diag}(d_1, \dots, d_p)$ for $d_1 \geq \dots \geq d_p$, and U_1 and V_1 are the first column vectors of U and V respectively. It is desirable to directly find the confidence interval for Λ_1 instead of $\tilde{\Lambda}_1$, however, as B is unobservable, $U_1 V_1^T$ is the “best guess” of the unit vector associated with Λ_1 in its direction.

To discuss the procedure in detail, in the matrix denoising problem of (2.5), the result from Taylor, Loftus and Tibshirani (2013) yields an exact conditional survival probability of d_1 as follows:

$$(2.7) \quad \mathbb{S}_{1, \tilde{\Lambda}_1} = \frac{\int_{d_1}^{\infty} e^{-\frac{(z-\tilde{\Lambda}_1)^2}{2\sigma^2}} z^{N-p} \prod_{j=2}^p (z^2 - d_j^2) dz}{\int_{d_2}^{\infty} e^{-\frac{(z-\tilde{\Lambda}_1)^2}{2\sigma^2}} z^{N-p} \prod_{j=2}^p (z^2 - d_j^2) dz} \sim \text{Unif}(0, 1).$$

When data is generated under $\Lambda_1 = 0$, then $\tilde{\Lambda}_1 = 0$ holds. In this case, (2.7) is the same as the test statistic $\mathbb{S}_{1,0}$ for testing $H_0 : \Lambda_1 = 0$ in Section 2.1.1, and yields exact type 1 error controls due to its uniformity. The proposed procedure for constructing the level α confidence interval is as follows:

$$(2.8) \quad CI = \{\delta : \min(\mathbb{S}_{1,\delta}, 1 - \mathbb{S}_{1,\delta}) > \alpha/2\}.$$

Since $\mathbb{S}_{1,\tilde{\Lambda}_1}$ in (2.7) is uniformly distributed, we observe that

$$\mathbb{P}(\tilde{\Lambda}_1 \in CI) = 1 - \alpha,$$

and thus (2.8) generates an exact level α confidence interval.

2.2. General hypothesis testing. In this section, we extend the test for the global null in Section 2.1.1 to a general test which investigates whether there exists the k th largest signal in B .

Suppose that we want to test the hypothesis

$$(2.9) \quad \begin{aligned} &H_{0,k} : \Lambda_k = 0 \text{ versus } H_{1,k} : \Lambda_k > 0 \\ \Leftrightarrow &H_{0,k} : \text{rank}(B) < k \text{ versus } H_{1,k} : \text{rank}(B) \geq k, \end{aligned}$$

for $k = 1, \dots, p - 1$. For $k = 1$, the null hypothesis in (2.9) corresponds to a global null as in Section 2.1.1. At $k = p$ the signal matrix B is full rank under the alternative hypothesis with $\Lambda_p \neq 0$. In this scenario, the problem of $\text{rank}(B) = p$ with the noise level of σ^2 , becomes unidentifiable with the problem of $\text{rank}(B) = p - 1$ with noise level of $\sigma^2 + \Lambda_p^2$. Thus, we assume $\text{rank}(B) < p$, and do not consider the case of $k = p$.

One of the most straightforward approaches for extending the global test (2.6) to testing (2.9) for $k = 1, \dots, p - 1$ would be to apply it sequentially, with removing the first observed $k - 1$ singular values at the k th step, and start at the k th observed singular value. That is, at the k th step, we can ignore the first $k - 1$ observed singular values of Y , and apply the test with plugging in the k th value to the place of the 1st value, the $(k + 1)$ th to the place of the 2nd value, and so on. This approach of ignoring the existence of the first $k - 1$ singular values is analogous to other methods dealing with essentially the same hypothesis testing [Muirhead (1982), Kritchman and Nadler (2008)].

How well does this work? Figure 2 shows an example. Here, $N = 20$, $p = 10$ and there is a rank one signal of moderate size. The top panels show quantile-quantile plots of the p -values for the aforementioned sequential *Kac-Rice test* versus the uniform distribution. We see that the p -values are small when the alternative hypothesis $H_{1,1}$ is true (step 1), and then fairly uniform for testing $\text{rank} \leq 1$ versus > 1 (step 2) as desired, which is the first case in which the null hypothesis is true. However, the test becomes more and more conservative for higher steps, although the p -values are generated under the null distributions. The conservativeness of these p -values can lead to potential loss of power. One of the reasons for

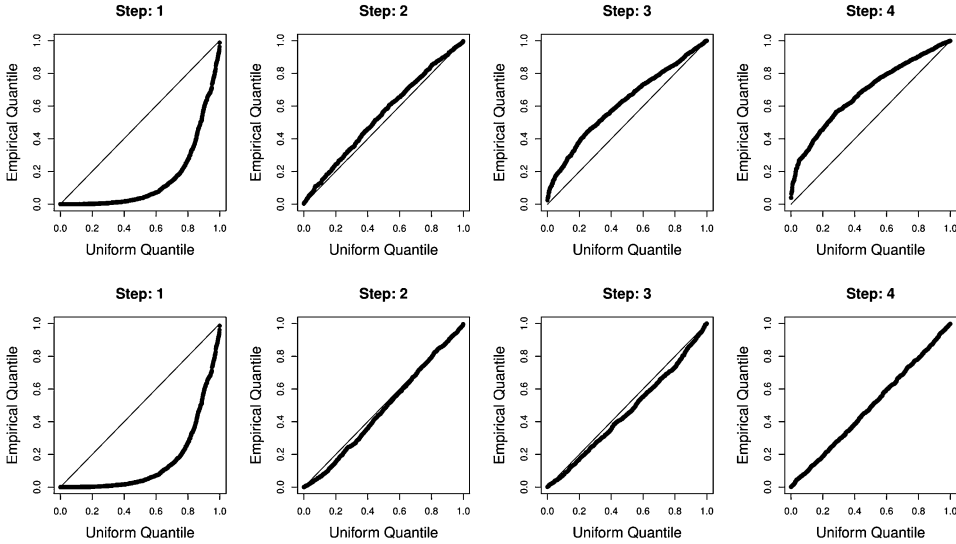


FIG. 2. *Quantile-quantile plots of the observed quantiles of p -values versus the uniform quantiles. With $N = 20$ and $p = 10$, the true rank of B is $\text{rank}(B) = 1$. The top panels are from the sequential Kac–Rice test and the bottom panels are from the CSV. The k th column represents quantile-quantile plots of the p -values for testing $H_{0,k} : \text{rank}(B) \leq k - 1$ for steps $k = 1, 2, 3, 4$.*

this conservativeness is that, at step $k = 3$, for example, the test does not consider that the two largest singular values have been removed. The test instead ignores the existence of the 1st and the 2nd singular values, and treats the 3rd one as the 1st singular value, plugging it into the place of the 1st singular value. As the 1st and the 3rd largest singular values do not have the same distribution with the density of the 1st singular value having more weights on large values, the sequential *Kac–Rice test* at $k = 3$ is no longer uniformly distributed; the test results in conservative p -values.

The plots in the bottom panel come from our proposed conditional singular value (CSV) test, to be described in Section 2.2.1. It follows the uniform distribution quite well for all null steps. For testing $H_{0,k} : \Lambda_k = 0$ at the k th step, the CSV method takes it into account that our interest is the k th signal by conditioning on the first $k - 1$ singular values when deriving its test statistic.

Section 2.2.1 discusses the CSV test procedure in detail. In Section 2.2.2, we propose the ICSV test, an integrated version of the CSV which has better power. Simulation results of the proposed procedures are illustrated in Section 2.2.3.

2.2.1. *The conditional singular value test.* In this section, we introduce a test in which the test statistic has an “almost exact” null distribution under $H_{0,k}$ in (2.9) for $k \in \{1, \dots, p - 1\}$.

Here are some notation used throughout the paper. We write the singular value decomposition of a signal matrix B as $B = U_B D_B V_B^T$, the singular value de-

composition of an observed data matrix Y as $Y = UDV^T$, and an $N \times r$ and an $N \times (p - r)$ column-wise submatrices of an $N \times p$ matrix M by $M_{[r]}$ and $M_{[-r]}$ where $M = [M_{[r]}|M_{[-r]}]$. Also, P_Q denotes a projection matrix onto a column-space of an $n_1 \times n_2$ matrix Q , and P_Q^\perp denotes a projection matrix onto a kernel of P_Q . that is, assuming $n_1 \geq n_2$ and Q is of full-rank, we have

$$P_Q = Q(Q^T Q)^{-1} Q^T,$$

$$P_Q^\perp = I_{n_1} - P_Q.$$

With these notation, the hypothesis in (2.9) can be rewritten as

$$(2.10) \quad \begin{aligned} H_{0,k} &: P_{U_{B[k-1]}}^\perp B P_{V_{B[k-1]}}^\perp = 0_{N \times p} \quad \text{versus} \\ H_{1,k} &: P_{U_{B[k-1]}}^\perp B P_{V_{B[k-1]}}^\perp \neq 0_{N \times p} \end{aligned}$$

since $P_{U_{B[k-1]}}^\perp B P_{V_{B[k-1]}}^\perp = 0_{N \times p}$ is equivalent to $\Lambda_k = \dots = \Lambda_p = 0$. The hypothesis in (2.10) examines whether the column spaces of $U_{B[k-1]}$ and $V_{B[k-1]}$ capture all nontrivial signals in B , or equivalently, whether the deflated residual space $(U_{B[-(k-1)]}, V_{B[-(k-1)]})$ contains any signals.

The proposed test procedure which we refer to as the conditional singular value test (CSV) is as follows:

TEST 2.1 (Conditional Singular Value test). With a given level α , and the following test statistic:

$$(2.11) \quad \mathbb{S}_{k,0} = \frac{\int_{d_k}^{d_{k-1}} e^{-\frac{z^2}{2\sigma^2}} z^{N-p} \prod_{j \neq k}^p |z^2 - d_j^2| dz}{\int_{d_{k+1}}^{d_{k-1}} e^{-\frac{z^2}{2\sigma^2}} z^{N-p} \prod_{j \neq k}^p |z^2 - d_j^2| dz},$$

where $d_0 = \infty$, we reject $H_{0,k}$ if $\mathbb{S}_{k,0} \leq \alpha$ and accept $H_{0,k}$ otherwise.

Analogous to (2.6), $\mathbb{S}_{k,0}$ compares the relative size of d_k ranging between (d_{k+1}, d_{k-1}) , and a small value of $\mathbb{S}_{k,0}$ implies a large value of d_k , supporting the alternative hypothesis $H_{1,k} : \Lambda_k > 0$. Likewise, the test statistic $\mathbb{S}_{k,0}$ plays the role of a p -value: it is the probability of observing larger values of the k th singular value than the actually observed d_k , under the distribution that is close to the null scenario. Here, $\mathbb{S}_{k,0}$ represents a survival probability of the k th singular value conditional on the observed $U_{[k-1]}, V_{[k-1]}$ and all the other singular values. As aforementioned, the distribution of $\mathbb{S}_{k,0}$ well represents the null. It is the conditional distribution of the k th singular value when the following situation holds:

$$(2.12) \quad P_{U_{[k-1]}}^\perp B P_{V_{[k-1]}}^\perp = 0_{N \times p}.$$

Here, (2.12) addresses the situation that $(U_{[k-1]}, V_{[k-1]})$ captures all the signals in B . Though (2.12) is close to $H_{0,k}$ in (2.10), there exists slight discrepancy between these two scenarios. As Y is drawn from a noisy matrix of B ,

$(U_{[k-1]}, V_{[k-1]})$ is also perturbed and does not coincide with $(U_{B[k-1]}, V_{B[k-1]})$. As a result, even though $H_{0,k} : \text{rank}(B) < k$ is true, the $(U_{[k-1]}, V_{[k-1]})$ cannot capture the entire signals and the bits of left-over remains in the deflated residuals. This happens especially when $k - 1 = \text{rank}(B)$. However, as the signals grow stronger and the step k advances further, the singular values $(U_{[k-1]}, V_{[k-1]})$ capture most of the signals, and thus the discrepancy between $H_{0,k}$ in (2.10) and (2.12) becomes slimmer.

Here, the data singular vectors $(U_{[k-1]}, V_{[k-1]})$ involved in (2.12) can be viewed as a component of a data-driven model of B with its rank being $k - 1$. As k goes beyond $\text{rank}(B)$, there is little information remaining to estimate the later singular vectors of B as their corresponding population singular value is 0. Nevertheless, the validity of the test statistic $\mathbb{S}_{k,0}$ depends only on finding candidate subspaces that contain the true singular vector subspaces of B . Continuing well beyond $\text{rank}(B)$ we will have added several left and right singular vectors that have little to do with B [cf. Theorem 4 of Paul (2007)]. The fact that (2.12) remains valid even well beyond $\text{rank}(B)$ at least partially explains why our tests continue to have roughly exact size while pseudo-rank becomes more and more conservative, as seen in Figure 3. Also, the bottom panels of Figure 2 confirm the claimed exact type I error control property of the procedure as after the true rank of one, the p -values are all close to uniform. Empirical results show close to exact type 1 error controls regardless of the magnitude of the signals even at $k - 1 = \text{rank}(B)$. Theorem 2.1 shows that the test statistic $\mathbb{S}_{k,0}$ is uniformly distributed when (2.12) holds, and thus yields close to exact type 1 error controls.

THEOREM 2.1.

If Y is drawn from $N(B, I_N \otimes I_p)$ and $P_{U_{[k-1]}}^\perp B P_{V_{[k-1]}}^\perp = 0_{N \times p}$ hold, then $\mathbb{S}_{k,0} \sim \text{Unif}(0, 1)$.

In constructing the test statistic $\mathbb{S}_{k,0}$, conditioning on $(U_{[k-1]}, V_{[k-1]})$ and (d_1, \dots, d_{k-1}) represents the selection event of choosing $k - 1$ active principal components at step k , and conditioning additionally on d_{k+1}, \dots, d_p leads to an inference of a saturated model. Though the test statistic $\mathbb{S}_{k,0}$ and its associated conditional density do not involve $(U_{[k-1]}, V_{[k-1]})$ since the singular vectors are independent from the singular values, it is incorporated in the test procedure as a component of the selection event of choosing $k - 1$ active principal components. By accounting for the selection event of active principal components along with their associated singular vectors, the proposed test procedure achieves (approximately) exact type 1 error controls; by constraining the test space to a deflated residual space via conditioning on the selection event, we can achieve the null distribution reflecting the current stage led by the selection procedure. This framework involves the post-selection inference as in the work of Tibshirani et al. (2014).

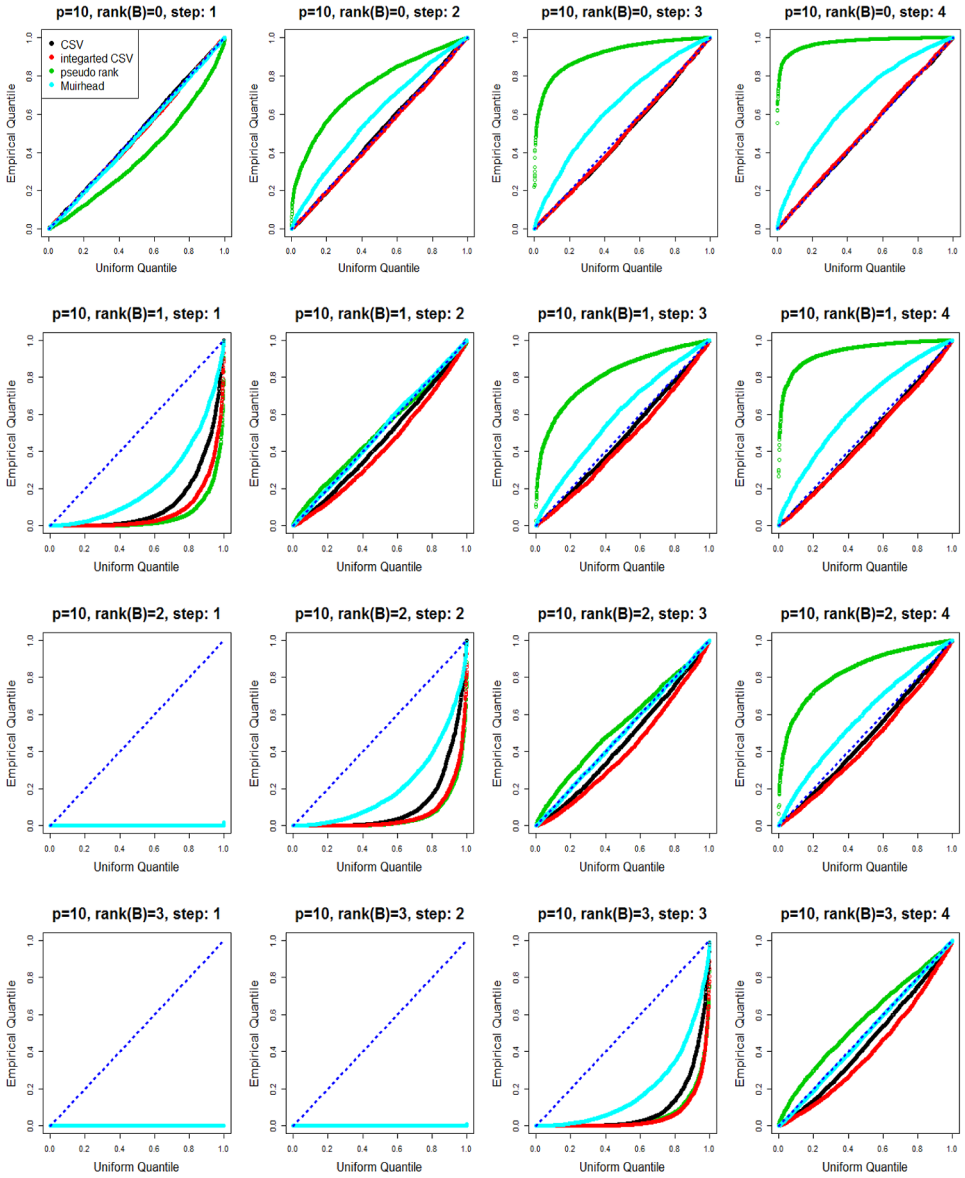


FIG. 3. *Quantile-quantile plots of the empirical quantiles of p -values versus the uniform quantiles when $p = 10$ and $N = 50$ at $m = 1.5$. From the top to the bottom, each row represents the case of the true rank(B) from 0 to 3. The columns represent the results for testing $H_{0,1}$ to $H_{0,4}$ from the left to the right. The six plots including (2nd row, 1st column), (3rd row, 1st column), (3rd row, 2nd column), (4th row, 1st column), (4th row, 2nd column) and (4th row, 3rd column) represent the cases under the alternative. The rest of the plots are under the null.*

The resulting conditional density of the singular values used in the test statistic reflects the fact that $(U_{[k-1]}, V_{[k-1]})$ is a sufficient statistic for B under (2.12). The detailed ideas of the test procedure and the proofs of the results are given in the supplementary material [Choi, Taylor and Tibshirani (2017)].

2.2.2. *The integrated conditional singular value test.* As a potential improvement of the CSV, we introduce an integrated version of $\mathbb{S}_{k,0}$ which we refer as Integrated Conditional Singular Value (ICSV) test. Our aim is to achieve higher power in detecting signals in B compared to the ordinary CSV. The idea is that conditioning on less can lead to greater power. The ordinary CSV test statistic $\mathbb{S}_{k,0}$ assumes a saturated model, conditioning on all the observed singular values except for the k th one. Here, we condition on only the first $k - 1$ observed singular values, and integrate out the last $p - k$ singular values with respect to the null distribution conditional on the first $k - 1$ observed singular values. This can be considered as averaging the last $p - k$ singular values across all the possible values of those with proper weights where the proper weights correspond to the conditional null distribution of the last $p - k$ singular values. The resulting statistic becomes a function of d_1, \dots, d_k , only the first k observed singular values of Y , which are associated with the selection event of active principal components that lead to step k . While the ordinary CSV test statistic $\mathbb{S}_{k,0}$ utilizes a saturated model by conditioning on all the observed singular values except for the k th one, the ICSV conditions only on the selection event and can be viewed as an inference on nonsaturated model. In its construction, the ICSV test utilizes a post-selection inference framework as in the work of Tibshirani et al. (2014) in the same manner as the ordinary CSV test.

The proposed test statistic is as follows:

$$(2.13) \quad \mathbb{V}_{k,0} = \frac{\int_{d_k}^{d_{k-1}} g(y_k; d_1, \dots, d_{k-1}) dy_k}{\int_0^{d_{k-1}} g(y_k; d_1, \dots, d_{k-1}) dy_k},$$

where

$$(2.14) \quad \begin{aligned} &g(y_k; d_1, \dots, d_{k-1}) \\ &= \int \dots \int \prod_{i=k}^p \left(e^{-\frac{y_i^2}{2\sigma^2}} y_i^{N-p} \right) \left(\prod_{i=k}^p \prod_{j>i} (y_i^2 - y_j^2) \right) \\ &\quad \cdot \left(\prod_{i=1}^{k-1} \prod_{j=k}^p (d_i^2 - y_j^2) \right) \mathbf{1}_{\{0 \leq y_p \leq y_{p-1} \leq \dots \leq y_k \leq d_{k-1}\}} dy_{k+1} \dots dy_p. \end{aligned}$$

Our proposed Integrated Conditional Singular Value (ICSV) test is as follows:

TEST 2.2 (ICSV test). With a given level α , we reject $H_{0,k}$ if $\mathbb{V}_{k,0} \leq \alpha$ and accept $H_{0,k}$ otherwise, where $\mathbb{V}_{k,0}$ is as defined in (2.13).

As in the CSV test, $\mathbb{V}_{k,0}$ performs as a p -value for the test. It examines the relative size of a gap from d_k to d_{k-1} , and a small value of $\mathbb{V}_{k,0}$ implies a large value of d_k , supporting the alternative hypothesis $H_{1,k} : \Lambda_k > 0$. In the ICSV test, the range of d_k is enlarged from (d_{k+1}, d_{k-1}) to $(0, d_{k-1})$ compared to the CSV test. The test statistic $\mathbb{V}_{k,0}$ is a survival probability of the k th singular value conditioned on the observed $U_{[k-1]}, V_{[k-1]}$ and the first $k - 1$ singular values with the distribution under (2.12), where the conditions reflect the data-driven selection event from the previous steps. In accordance with the CSV test, the test statistic of the ICSV test is uniformly distributed under (2.12), and thus provides close to exact type 1 error controls, which is shown in Theorem 2.2. Here, fixing the observed $(U_{[k-1]}, V_{[k-1]}, d_1, \dots, d_{k-1})$ corresponds to the selection of a data-driven model, and the proposed procedure investigates the deflated residual space resulting from conditioning on the selection event.

THEOREM 2.2.

If Y is drawn from $N(B, I_N \otimes I_p)$ and $P_{U_{[k-1]}}^\perp B P_{V_{[k-1]}}^\perp = 0_{N \times p}$ hold, then $\mathbb{V}_{k,0} \sim \text{Unif}(0, 1)$

Along with controlling type 1 errors to exact target levels, the ICSV test yields independent p -values when p -values are generated under (2.12). This independence relation between p -values is shown in Theorem 2.3, which corresponds to a special case of Theorem 4 by Fithian et al. (2015), page 15. The independence between p -values under the null hypothesis is a sufficient condition for a number of multiple hypothesis testing correction procedures [Benjamini and Hochberg (1995), G’Sell et al. (2013), Simes (1986)].

THEOREM 2.3.

If Y is drawn from $N(B, I_N \otimes I_p)$ and $P_{U_{[k-1]}}^\perp B P_{V_{[k-1]}}^\perp = 0_{N \times p}$ hold, then $\{\mathbb{V}_{i,0} | i \in \{k, \dots, p - 1\}\}$ are independent to each other.

Figure 3 demonstrates that the ICSV procedure achieves higher power than the ordinary CSV, and controls type 1 error for testing $H_{0,k} : \text{rank}(B) \leq k - 1$ near an exact target level as desired.

In this paper, we use importance sampling to evaluate the integral in (2.13) with samples drawn from the eigenvalues of a $(N - k + 1) \times (p - k + 1)$ Wishart matrix. As the computational cost increases sharply with large p , we are currently unable to compute this test for p beyond say 30 or 40. An interesting open problem is the numerical approximation of this integral, in order to scale the test to larger problems. We leave this as future work.

2.2.3. *Simulation examples.* In this section, we present results of the CSV and the ICSV for testing the general hypothesis $H_{0,k} : \text{rank}(B) \leq k - 1$ in (2.9) on simulated examples. We compare the performance of these proposed methods with those in [Kritchman and Nadler \(2008\)](#) and [Muirhead \(1982\)](#), Theorem 9.6.2 on page 409, mentioned in Section 1.2, which we refer as the *pseudorank* and the *Muirhead’s method*, respectively:

TEST 2.3 (Pseudorank). With a given level α , and following $\mu_{N,p}$ and $\sigma_{N,p}$,

$$\mu_{N,p} = \left(\sqrt{N - \frac{1}{2}} + \sqrt{p - \frac{1}{2}} \right)^2,$$

$$\sigma_{N,p} = \left(\sqrt{N - \frac{1}{2}} + \sqrt{p - \frac{1}{2}} \right) \left(\frac{1}{\sqrt{N - 1/2}} + \frac{1}{\sqrt{p - 1/2}} \right)^{1/3},$$

we reject $H_{0,k} : \text{rank}(B) \leq k - 1$ if

$$\frac{d_k^2 - \mu_{N,p-k}}{\sigma_{N,p-k}} > s(\alpha),$$

where $s(\alpha)$ is the upper α -quantile of the Tracy–Widom distribution.

TEST 2.4 (Muirhead’s method). With a given level α , and V_k defined as

$$V_k = \frac{(N - 1)^{q-1} \prod_{i=k}^p d_i^2}{\left(\frac{1}{q} \sum_{i=k}^p d_i^2 \right)^q},$$

we reject $H_{0,k} : \text{rank}(B) \leq k - 1$ if

$$-\left(N - k - \frac{2q^2 + q + 2}{6q} + \sum_{i=1}^{k-1} \frac{\bar{l}_q^2}{(d_i^2 - \bar{l}_q)^2} \right) \log V_k > \chi_{(q+2)(q-1)/2}^2(\alpha),$$

where $q = p - k + 1$, $\bar{l}_q = \sum_{i=k}^p d_i^2 / q$ and $\chi_m^2(\alpha)$ denotes the upper α quantile of the χ^2 distribution with degree m .

We investigate cases with i.i.d. Gaussian noise entries with $\sigma^2 = 1$. An $N \times p$ data matrix Y has the signal matrix B formed as follows:

$$(2.15) \quad B = U_B D_B V_B^T, \quad \Lambda_i = m \cdot i \cdot \sigma \sqrt{Np} \cdot I_{\{i \leq \text{rank}(B)\}},$$

where $D_B = \text{diag}(\Lambda_1, \dots, \Lambda_p)$ with $\Lambda_1 \geq \dots \geq \Lambda_p$, and U_B, V_B are rotation operators generated from a singular value decomposition of an $N \times p$ random Gaussian matrix with i.i.d. entries. The signals of B increase linearly. The constant m determines the magnitude of the signals. From $m = 1$, a *phase transition* phenomenon is observed when $\text{rank}(B) = 1$ in which the expectation of the largest singular value of Y starts to reflect the signal [[Nadler \(2008\)](#)].

We illustrate two cases of $(N, p) = (50, 10)$ and $(N, p) = (120, 100)$. For the former case, we set $m = 1.5$, with $\text{rank}(B) = 0, 1, 2, 3$. For the latter case, we set $m = 1.3$ with $\text{rank}(B) = 3, 5, 10, 50$. For both cases, we evaluate the procedure across 1000 repetition. The true known value of the noise level $\sigma^2 = 1$ is used for all testing procedures.

Figures 3 and 4 present quantile-quantile plots of the expected (uniform) quantiles versus the observed quantiles of p -values for testing the general hypothesis in (2.9). Under $H_{1,k}$, the ICSV test shows improved power compared to the CSV, and close to that of the *pseudorank* in Figure 3 [the six plots including (2nd row, 1st column), (3rd row, 1st column), (3rd row, 2nd column), (4th row, 1st column), (4th row, 2nd column) and (4th row, 3rd column)]. The first column of Figure 4 is under $H_{1,k}$ and illustrates that the CSV shows stronger power relative to the other methods as the true rank of B increases. For the instance of $N = 120$ and $p = 100$ with small true rank of B , the *Muirhead's* test shows anti-conservative performance under the null at the early steps which increases the risk of false discovery. Under $H_{0,k}$, both the CSV and the ICSV quantiles nearly agree with the expected quantiles and provide almost exact type 1 error controls as the theory predicts for all null steps. The results from both the *Pseudorank* and the *Muirhead's* test become strongly conservative for further steps.

2.3. Confidence interval construction. Here, we generalize the exact confidence interval construction procedure of the largest singular value in (2.8) to the k th signal parameter for any $k = 1, \dots, p - 1$.

We define the k th signal parameter $\tilde{\Lambda}_k$ as follows:

$$(2.16) \quad \tilde{\Lambda}_k = \langle U_k V_k^T, B \rangle,$$

where U_k and V_k are the k th column vector of U and V , respectively. We propose an approach to construct an exact level α confidence interval of $\tilde{\Lambda}_k$. Our proposed procedure is as follows:

$$(2.17) \quad CI_k(\mathbb{S}) = \{ \delta : \min(\mathbb{S}_{k,\delta}, 1 - \mathbb{S}_{k,\delta}) > \alpha/2 \},$$

where

$$\mathbb{S}_{k,\delta} = \frac{\int_{d_k}^{d_{k-1}} e^{-\frac{(z-\delta)^2}{2\sigma^2}} z^{N-p} \prod_{j \neq k}^p |z^2 - d_j^2| dz}{\int_{d_{k+1}}^{d_{k-1}} e^{-\frac{(z-\delta)^2}{2\sigma^2}} z^{N-p} \prod_{j \neq k}^p |z^2 - d_j^2| dz}.$$

We can find the boundary points of $CI_k(\mathbb{S})$ using bisection. Theorem 2.4 below shows that $CI_k(\mathbb{S})$ is an exact level α confidence interval. The proposed confidence interval construction procedure in (2.17) can be viewed as a variation of the previously proposed CSV test. The CSV test fixes $\delta = 0$, and yields a corresponding p -value. On the other hand, the confidence interval approach (2.17) provides pivots δ s that generate p -values within the target range.

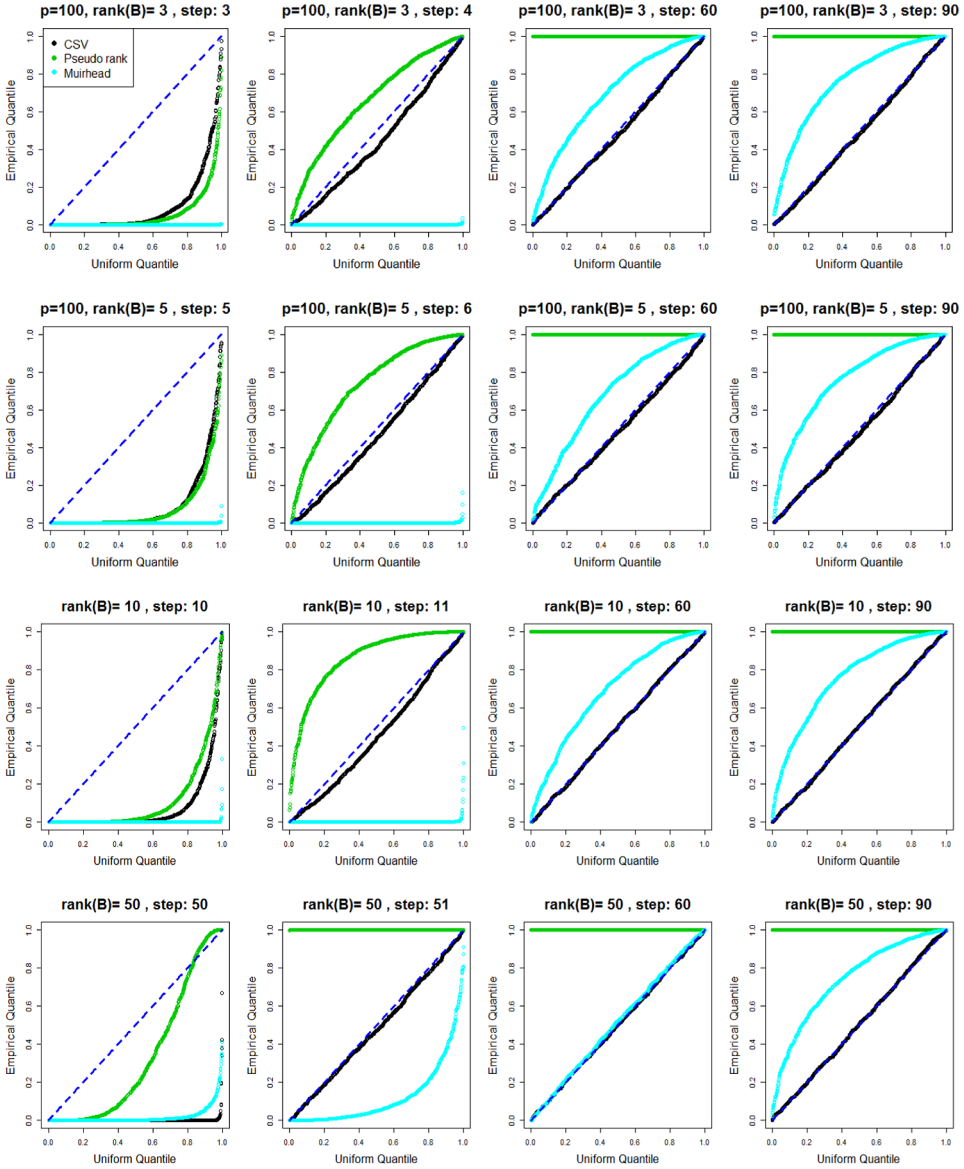


FIG. 4. *Quantile-quantile plots of the empirical quantiles of p-values versus the uniform quantiles when $p = 100$ and $N = 120$ at $m = 1.3$. From the top to the bottom, each row represents the case of κ , the true rank of B , equals 3, 5, 10 and 50. The columns represent the results for testing $H_{0,\kappa}$, $H_{0,\kappa+1}$, $H_{0,60}$ and $H_{0,90}$ from the left to the right. The first column represents the instances under the alternative. The rest of the columns are under the null.*

THEOREM 2.4. $\mathbb{S}_{k,\delta}$ is uniformly distributed when $\delta = \tilde{\Lambda}_k$.

Figure 5 shows the coverage rates of $CI_k(\mathbb{S})$. As expected, the coverage rate of the true parameter is close to the target $1 - \alpha$.

2.3.1. *Simulation studies of the confidence interval construction.* We illustrate the coverage rates of $CI_k(\mathbb{S})$ in (2.17) on simulated data. The simulation settings are the same as in Section 2.2.3 with $N = 50$ and $p = 10$. Figure 5 shows the coverage rate of the 95% confidence intervals for the first two signal parameters $\tilde{\Lambda}_1$ and $\tilde{\Lambda}_2$ in (2.16). Here, we vary m in (2.15) from 0 to 2. Large m leads to large magnitude of true $\tilde{\Lambda}_1$ and $\tilde{\Lambda}_2$.

Figure 5 shows that regardless of the step $k = 1, 2$, the true value of $\text{rank}(B)$, or the size of $\tilde{\Lambda}_k$, the constructed confidence intervals cover the parameters at the desired level as expected.

3. Rank estimation. This section discusses the selection of the number of principal components, or equivalently, the estimation of the true rank of B under our model assumptions. For determining the rank of B , here we investigate the *StrongStop* procedure [G’Sell et al. (2013)], applying it to the tests developed in Section 2.2 (CSV and ICSV).

We determine the rank of B based on our testing results on $H_{0,k} : \text{rank}(B) \leq k - 1$ in (2.9), since $H_{0,k}$ explicitly tests the range of the true rank. Given the sequence of hypothesis $H_{0,k} : \text{rank}(B) \leq k - 1$ with $k = 1, \dots, p - 1$, the rejection of these must be carried out in a sequential fashion such that once $H_{k,0}$ is rejected, all $H_{\gamma,0}$ for $\gamma \leq k$ should be rejected as well. Under such sequential testing framework of this kind, it is natural to choose $\text{rank}(B)$ to be the largest k that rejects $H_{0,k}$. The question here is how to choose the ‘stopping point’ for rejection.

One of the simplest methods is to choose the value k at which $H_{0,k}$ is rejected for the last time with a given level α :

$$\hat{k}_{\text{simple}} = \max\{k \in \{1, \dots, p - 1\} : p_k \leq \alpha\},$$

which we refer as *SimpleStop*.

In this paper, instead of the *SimpleStop*, we use the *StrongStop*. This procedure takes sequential p -values as its input and controls family-wise error rate. When the p -values under the null are drawn from uniform distribution, independent to each other, the *StrongStop* procedure controls the family-wise error rate under a given level of α [G’Sell et al. (2013)], Theorem 3. For rank determination, by the nature of our hypothesis, $H_{0,\gamma} : \text{rank}(B) \leq \gamma - 1$ being true implies $H_{0,k} : \text{rank}(B) \leq k - 1$ also being true for all $k > \gamma$. The family-wise error rate control property in rank determination, therefore, becomes control of rank over-estimation with level α as follows:

$$\mathbb{P}(\hat{k} > \kappa) \leq \alpha,$$

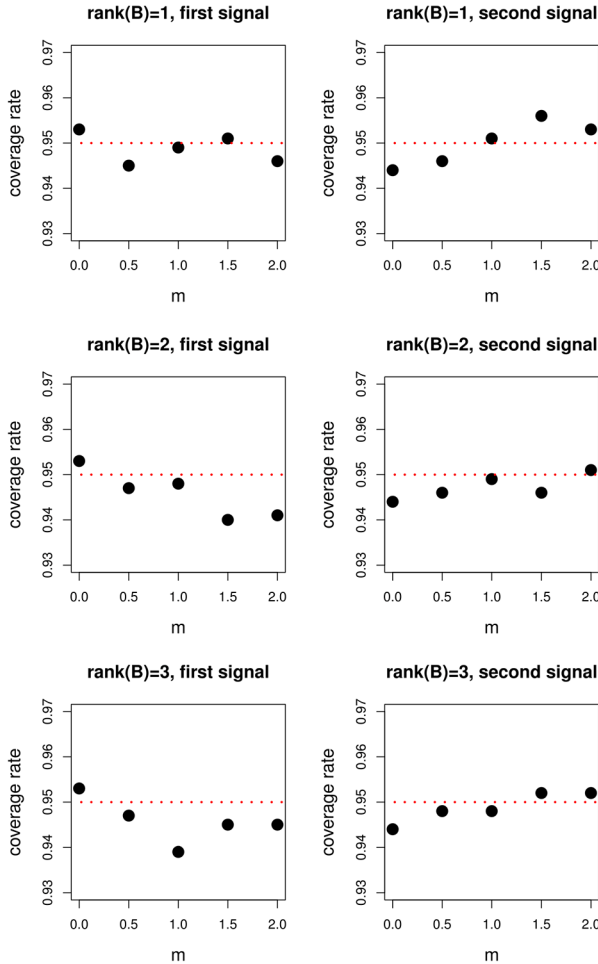


FIG. 5. Coverage rate of $\tilde{\Lambda}_1$ and $\tilde{\Lambda}_2$ versus the magnitude m in (2.15) with $N = 50$ and $p = 10$. The first and the second signal correspond to $\tilde{\Lambda}_1$ and $\tilde{\Lambda}_2$, respectively. Coverage rate denotes the proportion of times that the constructed confidence intervals from $CI_k(\mathbb{S})$ covered the true parameter. From the top to the bottom, each row represents the case of the true rank(B) being 1 to 3. The columns illustrate the results regarding $\tilde{\Lambda}_1$ and $\tilde{\Lambda}_2$ from the left to the right.

where $\hat{\kappa}$ denotes the selected rank(B) = κ . The proposed procedure is as follows:

$$\hat{\kappa} = \max \left\{ k \in \{1, \dots, p - 1\} : \exp \left(\sum_{j=k}^{p-1} \frac{\log p_j}{j} \right) \leq \frac{\alpha k}{p - 1} \right\}.$$

Here, p_k denotes the value of either $\mathbb{S}_{k,0}$ of the CSV or $\mathbb{V}_{k,0}$ of the ICSV, and conventionally $\max(\emptyset) = 0$. As shown in Theorem 2.2 and Theorem 2.3, under the general null steps in (2.9), the p -values from the ICSV test are almost uniformly and independently distributed to each other, practically satisfying the the

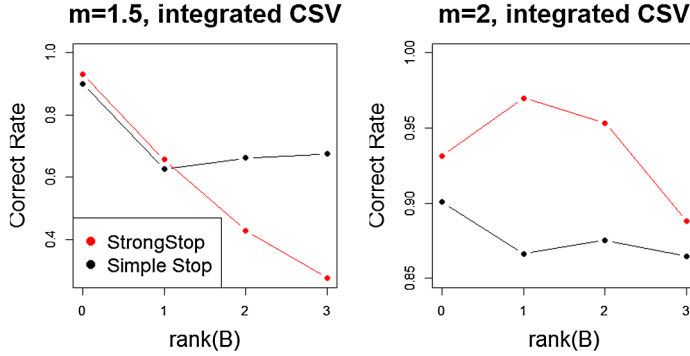


FIG. 6. Rate of selecting the correct rank(B) versus rank(B) for the ICSV procedure with $N = 50$ and $p = 10$. From the left to the right, m increases from 1.5 to 2 where m determines the magnitude of the signals in (2.15). The black dots and the red dots represent the results from the SimpleStop and the StrongStop respectively. For both stopping rules, $\alpha = 0.1$ is used.

assumptions of the *StronStop*. For the CSV test, the independence property is not guaranteed.

Figure 6 illustrates the simulation results of the *StrongStop* on p -values from the ICSV procedure with $\alpha = 0.1$. The simulation setting is the same as in Section 2.2.3 with $N = 50$ and $p = 10$. We compare the *StrongStop* with the *SimpleStop*. From Figure 6, we observe that for weaker signals, the *SimpleStop* tends to choose the correct rank of B more often than the *StrongStop* while for strong signals, the *StrongStop* shows better performance in general.

4. Estimating the noise level. For the testing procedures CSV and ICSV, and the confidence interval $CI_k(\mathbb{S})$ construction procedure, we have assumed that the noise level σ^2 is known. In case the prior information of σ^2 is unavailable, the value of σ^2 needs to be estimated. In this section, we introduce a data-driven method for estimating σ^2 .

For the estimation of σ^2 , it is popular to assume that the rank of B is known. One of the simplest methods estimates σ^2 using mean of sum of squared residuals by

$$\hat{\sigma}_{\text{simple}}^2 = \frac{1}{N(p - \kappa)} \sum_{j=\kappa+1}^p d_j^2$$

with known rank(B) = κ .

Instead of using the rank of B , Gavish and Donoho (2014) use the median of the singular values of Y as a robust estimator of σ^2 as follows:

$$(4.1) \quad \hat{\sigma}_{\text{med}}^2 = \frac{d_{\text{med}}^2}{N \cdot \mu_\beta},$$

where d_{med} is a median of the singular values of Y and μ_β is a median of a Marčenko-Pastur distribution with $\beta = N/p$. This estimator works under the assumption that $\text{rank}(B) \ll \min(N, p)$.

In this paper, we propose an estimator, which makes few assumptions. Our approach uses cross-validation and the sum of squared residuals as an extension of classical noise level estimator. For a fixed value of λ , we define our estimator $\hat{\sigma}_{\lambda,df,c}^2$ as follows:

$$(4.2) \quad \hat{\sigma}_{\lambda,df,c}^2 = \frac{1}{N \cdot (p - c \cdot df_{\hat{B}_\lambda})} \|Y - \hat{B}_\lambda\|_F^2 \quad \text{for } c \in [0, 1]$$

with

$$\hat{B}_\lambda = \underset{B \in \mathbb{R}^{N \times p}}{\text{argmin}} \frac{1}{2} \|Y - B\|_F^2 + \lambda \|B\|_*$$

$$df_{\hat{B}_\lambda} = \sum_{k=1}^p 1\{l_{\lambda,k} > 0\},$$

where $l_{\lambda,k}$ denotes the k th singular value of \hat{B}_λ . For $c = 0$, the estimator $\hat{\sigma}_{\lambda,df,c}^2$ represents the ordinary mean squared residual, and for $c = 1$, it accounts for the degrees of freedom as proposed by Reid, Tibshirani and Friedman (2016). With $c \in (0, 1)$, the performance of the estimator lies between $\hat{\sigma}_{\lambda,df,0}^2$ and $\hat{\sigma}_{\lambda,df,1}^2$. We use cross-validation for choosing the appropriate value of the regularization parameter λ .

For the estimator $\hat{\sigma}_{\lambda,df,c}^2$, choosing an appropriate value for the regularization parameter λ is important, since \hat{B}_λ depends on λ . In penalized regression, it is common to use cross-validation for this purpose, examining a grid of λ values. Unlike the regression setting, however, here there is no outcome variable, and thus it is not clear how to make predictions on left-out data points. In this paper, we use the *softImpute* algorithm [Mazumder, Hastie and Tibshirani (2010)] to address this issue. In the presence of missing values in a given data matrix, *softImpute* carries out matrix completion with the following criterion:

$$(4.3) \quad \min_B \frac{1}{2} \|P_\Omega(Y) - P_\Omega(B)\|_F^2 + \lambda \|B\|_*,$$

where Ω is an index set of observed data points with a function $P_\Omega(\cdot)$ such that $P_\Omega(Y)_{(i,j)} = Y_{i,j}$ if $(i, j) \in \Omega$ and 0 otherwise. We define the prediction error of the unobserved values in Ω as follows:

$$\text{err}_\lambda(\Omega) = \|P_{\tilde{\Omega}}(Y) - P_{\tilde{\Omega}}(\hat{B}_\lambda^S(\Omega))\|_F^2,$$

where $\hat{B}_\lambda^S(\Omega)$ denotes the estimator of B acquired from (4.3) and $\tilde{\Omega}$ denotes the index set of unobserved values in Ω ($\tilde{\Omega} = \Omega^c$). Using this prediction error, we carry out k -fold cross-validation, randomly generating k nonoverlapping leave-out

TABLE 1

Simulation results for estimating the noise level $\sigma^2 = 1$ with $N = 50$ and $p = 10$. We vary $\text{rank}(B)$ from 0 to 3, and m from 0.5 to 2.0. Each column represents the mean estimated value of σ^2 (“Est”), and standard error (“se”) of the corresponding estimator

m	$\hat{\sigma}_{\lambda_{CV},df,0}^2$		$\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$		$\hat{\sigma}_{\lambda_{CV},df,1}^2$		$\hat{\sigma}_{\text{med}}^2$	
	Est	se	Est	se	Est	se	Est	se
rank(B) = 0								
0.0	0.863	0.230	0.926	0.210	0.990	0.230	0.996	0.084
rank(B) = 1								
0.5	0.869	0.236	0.934	0.216	1.000	0.238	1.006	0.085
1.0	0.869	0.254	0.954	0.234	1.039	0.263	1.027	0.088
1.5	0.823	0.271	0.969	0.254	1.127	0.318	1.044	0.090
2.0	0.789	0.246	0.999	0.224	1.245	0.324	1.052	0.091
rank(B) = 2								
0.5	0.871	0.260	0.962	0.237	1.061	0.280	1.039	0.089
1.0	0.784	0.262	1.025	0.244	1.321	0.355	1.093	0.096
1.5	0.700	0.288	1.052	0.201	1.611	0.431	1.121	0.100
2.0	0.646	0.211	1.047	0.196	1.762	0.471	1.134	0.102
rank(B) = 3								
0.5	0.827	0.303	1.002	0.288	1.206	0.365	1.098	0.095
1.0	0.674	0.235	1.076	0.231	1.771	0.459	1.186	0.107
1.5	0.585	0.202	1.062	0.228	2.135	0.623	1.226	0.113
2.0	0.549	0.180	1.057	0.224	2.324	0.711	1.242	0.116

sets of size $\frac{N-p}{k}$ from Y . For a grid of λ values, we compute the average of err_λ for each λ over the k left-out data sets. We choose our λ to be the minimizer of the average err_λ as in usual cross-validation [see, e.g., [Hastie, Tibshirani and Friedman \(2009\)](#)].

4.1. *A study of noise level estimation.* We illustrate simulation examples of noise level estimation of the proposed method $\hat{\sigma}_{\lambda,df,c}^2$ with $c = 0, \frac{2}{3}$ and 1 in (4.2) compared to $\hat{\sigma}_{\text{med}}^2$ in (4.1). These approaches do not require predetermined knowledge of $\text{rank}(B)$.

Simulation settings are the same as in Section 2.2.3 with $N = 50$ and $p = 10$. The true value of the noise level is $\sigma^2 = 1$ and for choosing λ , 20-fold cross-validation is used. Table 1 illustrates the simulation results for the proposed estimators.

In this setting, $\hat{\sigma}_{\lambda_{CV},df,0}^2$ decreases with larger $\text{rank}(B)$ and signals while $\hat{\sigma}_{\lambda_{CV},df,1}^2$ and $\hat{\sigma}_{\text{med}}^2$ increases. For large $\text{rank}(B)$ and signals, $\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$ shows good results, as compared to other methods. The poor performance of $\hat{\sigma}_{\lambda_{CV},df,1}^2$ may be caused

by the use of an improper definition of df , the degrees of freedom. Following the definition of degrees of freedom by [Efron et al. \(2004\)](#), our simulation result shows that the number of nonzero singular values does not coincide with degrees of freedom under our setting. Further investigation into the degrees of freedom is needed in future work.

The competing method $\hat{\sigma}_{\text{med}}^2$ consistently shows small standard deviations. However, with large $\text{rank}(B)$, especially when $\text{rank}(B) \geq p/2$, the procedures over-estimates σ^2 due to the effect of the signals.

5. Additional examples. We discuss additional examples in this section. Section 5.1 presents results of the proposed methods when the estimated noise level is used. In Section 5.2, hypothesis testing results with non-Gaussian noise are illustrated. Section 5.3 shows results on some real data.

5.1. *Simulation examples with unknown noise level.* In this section, we illustrate the results when estimated σ^2 value is used on simulated data. For the estimation of the noise level, we use $\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$ and $\hat{\sigma}_{\text{med}}^2$ which showed good performance in Section 4.1. As in Section 4.1, for the estimator $\hat{\sigma}_{\lambda_{CV},df,c}^2$, 20-fold cross-validation and $c = 2/3$ is used. The simulation settings are the same as in Section 2.2.3 with $N = 50$ and $p = 10$. We investigate the case of $m = 1.5$.

Table 2 illustrates the estimated σ^2 values we used for the testing procedure. Figure 7 shows quantile-quantile plots of observed p -values obtained from using the estimated σ^2 versus the expected (uniform) quantiles. In quantile-quantile plots, both estimators of σ^2 show reasonable results in general, and for large $\text{rank}(B)$, $\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$ shows better result than $\hat{\sigma}_{\text{med}}^2$. In terms of coverage rate of confidence interval, we can see from Figure 8 that $\hat{\sigma}_{\text{med}}^2$ dominates for all cases, which might be due to small standard deviation of $\hat{\sigma}_{\text{med}}^2$ estimator. The estimation

TABLE 2
Simulation results for estimating the noise level $\sigma^2 = 1$ at $m = 1.5$ with $N = 50$ and $p = 10$. We vary the rank of B from 0 to 3. Shown are the mean estimated noise level (“Est”) and standard error (“se”) of the corresponding estimators

rank(B)	$\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$		$\hat{\sigma}_{\text{med}}^2$	
	Est	se	Est	se
0	0.926	0.210	0.996	0.084
1	0.969	0.254	1.044	0.090
2	1.052	0.201	1.121	0.100
3	1.062	0.228	1.226	0.113

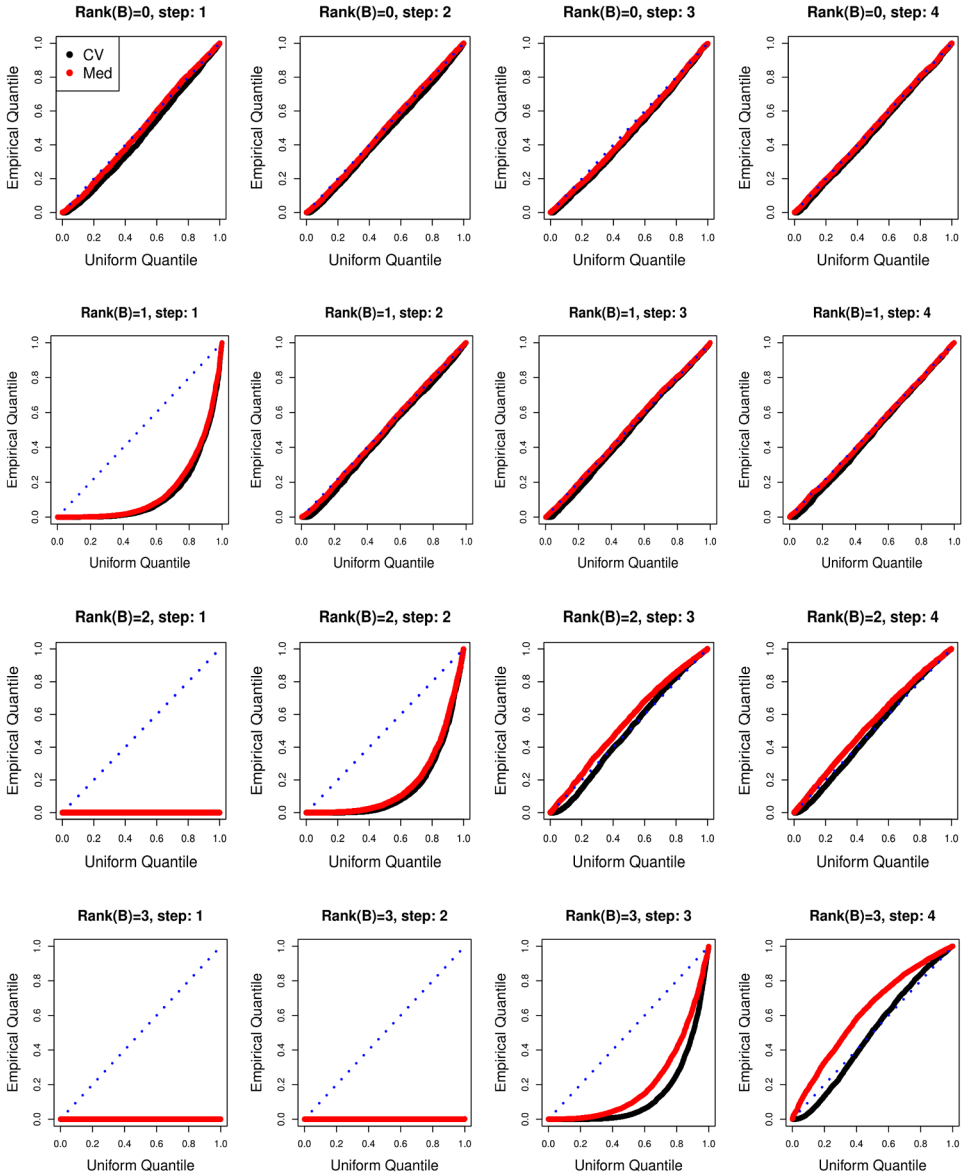


FIG. 7. *Quantile-quantile plots of the empirical quantiles of p -values with estimated σ^2 versus the uniform quantiles at $m = 1.5$ with $N = 50$ and $p = 10$. Empirical p -values are from the CSV test. Each row represents $\text{rank}(B) = 0$ to $\text{rank}(B) = 3$ from the top to the bottom. Each column represents the 1st test to the 4th test from the left to the right. The black dots (CV) and the red dots (med) represent p -values from using $\hat{\sigma}_{\lambda_{CV}, df, \frac{2}{3}}^2$ and $\hat{\sigma}_{\text{med}}^2$, respectively.*

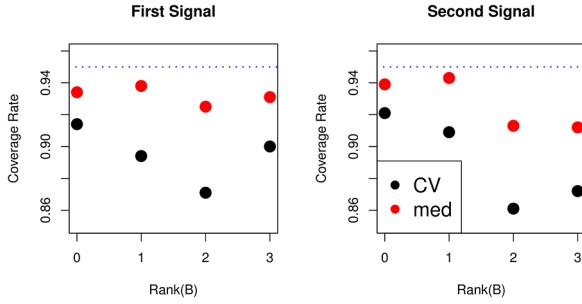


FIG. 8. Coverage rate versus rank(B) of the first two signal parameters at $m = 1.5$ using estimated σ^2 with $N = 50$ and $p = 10$. Coverage rate denotes proportion of times that the constructed confidence interval from $CI_k(S)$ covered the true parameter. The black dots (CV) and the red dots (med) represent the estimation using $\hat{\sigma}_{\lambda_{CV}, df, \frac{2}{3}}^2$ and $\hat{\sigma}_{med}^2$, respectively.

of rank(B) is presented in Table 3. For the estimation, the *StrongStop* is applied to the CSV p -values with level $\alpha = 0.05$. The estimation performance seems to vary with the quality of the estimate of σ^2 .

5.2. *Simulation example with non-Gaussian noise.* Our testing procedure is based on an assumption of Gaussian noise. Here, we investigate the performance of the CSV test on simulated examples with the normality assumption of noise violated. Simulation settings are the same as in Section 2.2.3 with $N = 50$ and $p = 10$ except for the noise distribution. We study the case of rank(B) = 1 with $m = 1.5$ along with two sorts of noise distribution: heavy tailed and right skewed. The heavy tailed noise is drawn from $\sqrt{\frac{3}{5}}t_5$ where t_5 denotes t -distribution with degrees of

TABLE 3
Simulation results for selecting an exact rank(B) from the CSV test using estimated σ^2 at $m = 1.5$ with $N = 50$ and $p = 10$. The *StrongStop* is applied to sequential p -values with $\alpha = 0.05$. We vary the rank of B from 0 to 3. Shown are the rate of selecting the correct rank of B (“*Rate*”) and mean squared error (“*MSE*”) using the corresponding estimator of σ^2

rank(B)	$\hat{\sigma}_{\lambda_{CV}, df, \frac{2}{3}}^2$		$\hat{\sigma}_{med}^2$	
	Rate	MSE	Rate	MSE
0	0.894	3.704	0.948	0.063
1	0.455	4.243	0.486	0.514
2	0.231	1.080	0.157	0.853
3	0.181	0.845	0.026	0.975

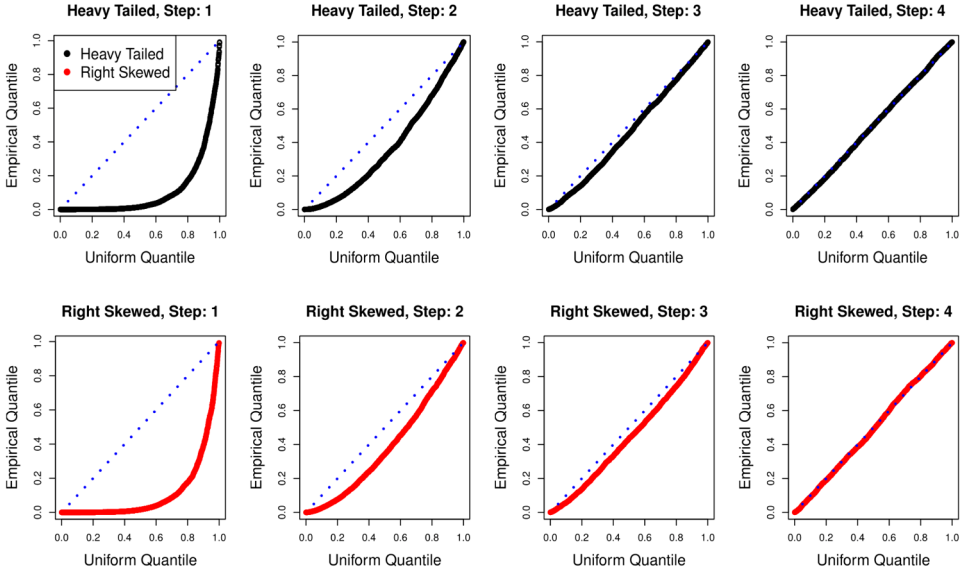


FIG. 9. *Quantile-quantile plots of empirical quantiles of p -values versus uniform quantiles when noises are drawn i.i.d. from the heavy tailed or the right skewed distribution with known value of $\sigma^2 = 1$ with $N = 50$ and $p = 10$. The true rank of B is $\text{rank}(B) = 1$ with $m = 1.5$. The top panels correspond to the heavy tailed noise from $\sqrt{\frac{3}{5}}t_5$ and the bottom panels correspond to the right skewed noise from $\sqrt{\frac{3}{10}}t_5 + \sqrt{\frac{1}{2}}(\exp(1) - 1)$. Each column represents the 1st test to the 4th test from the left to the right.*

freedom = 5, and the right skewed noises are drawn from $\sqrt{\frac{3}{10}}t_5 + \sqrt{\frac{1}{2}}(\exp(1) - 1)$ where $\exp(1)$ denotes the exponential distribution with mean = 1. In each case, noise entries are drawn i.i.d., and known value of $\sigma^2 = 1$ is used.

Figure 9 shows quantile-quantile plots of the observed p -values versus expected (uniform) quantiles. The top panels correspond to the heavy tailed noise from $\sqrt{\frac{3}{5}}t_5$ and the bottom panels correspond to the right skewed noise from $\sqrt{\frac{3}{10}}t_5 + \sqrt{\frac{1}{2}}(\exp(1) - 1)$. Quantile-quantile plots from both types of noise show that the p -values deviate slightly from a uniform distribution under the null $H_{0,2} : \text{rank}(B) \leq 1$. For further steps, p -values lie closer to the reference line.

The nonconformity shown in early steps under the null hypothesis is not surprising considering the construction of the CSV procedure based on Gaussian noise. In future work, we will investigate whether the procedures introduced here can be extended to a method robust to nonnormality using the data-oriented method such as bootstrap.

5.3. *Real data example.* In this section, we revisit the real data example mentioned in Figure 1. We apply the CSV test to the centered data of examination

TABLE 4

P-values at each step (“Step”) and the selected number of principal components (“Selected”) from the CSV test using the estimated noise level on the examination marks of 88 students on five different topics (Mechanics, Vectors, Algebra, Analysis and Statistics). The StrongStop is applied to select the number of principal components at level $\alpha = 0.05$. The estimated value of σ^2 is 37.560 for the estimator $\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$ and 104.551 for the estimator $\hat{\sigma}_{med}^2$

Step	$\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$	$\hat{\sigma}_{med}^2$
1	0.000	0.000
2	0.000	0.000
3	0.000	0.739
4	0.000	0.551

	$\hat{\sigma}_{\lambda_{CV},df,\frac{2}{3}}^2$	$\hat{\sigma}_{med}^2$
Selected	4	2

marks of 88 students on 5 different topics of Mechanics, Vectors, Algebra, Analysis and Statistics [Mardia, Kent and Bibby (1979), pages 3–4], and determine the number of principal components to retain for PCA.

In this data, Mechanics and Vectors were closed book exams while the other topics were open book exam. We use $\hat{\sigma}_{\lambda_{CV},df,c}^2 = 37.560$ and $\hat{\sigma}_{med}^2 = 104.551$ for the estimated noise level. For $\hat{\sigma}_{\lambda_{CV},df,c}^2$, 20-fold cross-validation is used with $c = 2/3$. The CSV test results are presented in Table 4. The estimated $\text{rank}(B)$ varies with the estimator of σ^2 in use: the estimated $\text{rank}(B)$ is 4 with $\hat{\sigma}_{\lambda_{CV},df,c}^2$ estimator and 2 with $\hat{\sigma}_{med}^2$ with level $\alpha = 0.05$ using the *StrongStop*. Thus, in PCA we may use two or four principal components depending on our choice of the noise level.

6. Conclusions. In this paper, we have proposed distribution-based methods for choosing the number of principal components of a data matrix. We have proposed novel methods both for hypothesis testing and the construction of confidence intervals of the signals. The methods have exact type I error control and show promising results in simulated examples. We have also introduced data-based methods for estimating the noise level.

There are many topics that deserve further investigation. In following studies, the analysis of power of the proposed tests and the width of the constructed confidence interval will be investigated. Also, application of the methods to high dimensional data using numerical approximations will be explored. For multiple

hypothesis testing corrections to be properly applied, we will study the dependence structure of the p -values between different steps. In addition, for robustness to non-Gaussian noise, bootstrap versions of this procedure will be investigated. Future work may involve a notion of degrees of freedom of the spectral estimator of the signal matrix. These extensions may lead to improvement in noise level estimation.

Variations of these procedures can potentially be applied to canonical correlation analysis (CCA) and linear discriminant analysis (LDA), and these are topics for future work.

Acknowledgements. We would like to thank Boaz Nadler and Iain Johnstone for helpful conversations.

SUPPLEMENTARY MATERIAL

Supplement to “Selecting the number of principal components: estimation of the true rank of a noisy matrix” (DOI: [10.1214/16-AOS1536SUPP](https://doi.org/10.1214/16-AOS1536SUPP); .pdf). The supplement to this paper [Choi, Taylor and Tibshirani (2017)] contains additional details of the proposed CSV and ICSV tests, and proofs of the main theorems.

REFERENCES

- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](#)
- CAI, J.-F., CANDÈS, E. J. and SHEN, Z. (2010). A singular value thresholding algorithm for matrix completion. *SIAM J. Optim.* **20** 1956–1982. [MR2600248](#)
- CHOI, Y., TAYLOR, J. and TIBSHIRANI, R. (2017). Supplement to “Selecting the number of principal components: Estimation of the true rank of a noisy matrix.” DOI:[10.1214/16-AOS1536SUPP](https://doi.org/10.1214/16-AOS1536SUPP).
- EFRON, B., HASTIE, T., JOHNSTONE, I. and TIBSHIRANI, R. (2004). Least angle regression. *Ann. Statist.* **32** 407–499. [MR2060166](#)
- FITHIAN, W., TAYLOR, J., TIBSHIRANI, R. and TIBSHIRANI, R. (2015). Selective Sequential Model Selection. Preprint. Available at [arXiv:1512.02565](https://arxiv.org/abs/1512.02565).
- G’SELL, M. G., WAGER, S., CHOULDECHOVA, A. and TIBSHIRANI, R. (2013). Sequential Selection Procedures And False Discovery Rate Control. Preprint. Available at [arXiv:1309.5352](https://arxiv.org/abs/1309.5352).
- GAVISH, M. and DONOHO, D. L. (2014). The optimal hard threshold for singular values is $4/\sqrt{3}$. *IEEE Trans. Inform. Theory* **60** 5040–5053. [MR3245370](#)
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. Springer, New York. [MR2722294](#)
- JAMES, A. T. (1964). Distributions of matrix variates and latent roots derived from normal samples. *Ann. Math. Stat.* **35** 475–501. [MR0181057](#)
- JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* **29** 295–327. [MR1863961](#)
- JOLLIFFE, I. T. (2002). *Principal Component Analysis*, 2nd ed. Springer, New York. [MR2036084](#)
- JOSSE, J. and HUSSON, F. (2012). Selecting the number of components in principal component analysis using cross-validation approximations. *Comput. Statist. Data Anal.* **56** 1869–1879. [MR2892383](#)
- KRITCHMAN, S. and NADLER, B. (2008). Determining the number of components in a factor model from limited noisy data. *Chemom. Intell. Lab. Syst.* **94** 19–32.

- MARDIA, K. V., KENT, J. T. and BIBBY, J. M. (1979). *Multivariate Analysis*. Academic Press [Harcourt Brace Jovanovich, Publishers], London-New York-Toronto, Ont. [MR0560319](#)
- MAZUMDER, R., HASTIE, T. and TIBSHIRANI, R. (2010). Spectral regularization algorithms for learning large incomplete matrices. *J. Mach. Learn. Res.* **11** 2287–2322.
- MUIRHEAD, R. J. (1982). *Aspects of Multivariate Statistical Theory*. Wiley, New York. [MR0652932](#)
- NADLER, B. (2008). Finite sample approximation results for principal component analysis: A matrix perturbation approach. *Ann. Statist.* **36** 2791–2817. [MR2485013](#)
- PAUL, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statist. Sinica* **77** 1617–1642.
- REID, S., TIBSHIRANI, R. and FRIEDMAN, J. (2016). A study of error variance estimation in Lasso regression. *Statist. Sinica* **26** 35–67. [MR3468344](#)
- SIMES, R. J. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika* **73** 751–754. [MR897872](#)
- TAYLOR, J., LOFTUS, J. and TIBSHIRANI, R. (2013). Tests in adaptive regression via the Kac-Rice formula. Preprint. Available at [arXiv:1308.3020](#).
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- TIBSHIRANI, R., TAYLOR, J., LOCKHART, R. and TIBSHIRANI, R. (2014). Exact Post-Selection Inference for Sequential Regression Procedures. Preprint. Available at [arXiv:1401.3889](#).

Y. CHOI
 DEPARTMENT OF STATISTICS
 AND APPLIED PROBABILITY
 NATIONAL UNIVERSITY OF SINGAPORE
 BLOCK S16, LEVEL 6
 SCIENCE DRIVE 2
 SINGAPORE 117546
 SINGAPORE
 E-MAIL: stachoiy@nus.edu.sg

J. TAYLOR
 DEPARTMENT OF STATISTICS
 STANFORD UNIVERSITY
 390 SERRA MALL
 STANFORD, CALIFORNIA 94305
 USA
 E-MAIL: jonathan.taylor@stanford.edu

R. TIBSHIRANI
 DEPARTMENT OF HEALTH, RESEARCH & POLICY
 DEPARTMENT OF STATISTICS
 390 SERRA MALL
 STANFORD UNIVERSITY
 STANFORD, CALIFORNIA 94305
 USA
 E-MAIL: tibs@stanford.edu