
Robust SVD Made Easy: A fast and reliable algorithm for large-scale data analysis

Sangil Han
Seoul National University

Kyoowon Kim
Seoul National University

Sungkyu Jung
Seoul National University

Abstract

The singular value decomposition (SVD) is a crucial tool in machine learning and statistical data analysis. However, it is highly susceptible to outliers in the data matrix. Existing robust SVD algorithms often sacrifice speed for robustness or fail in the presence of only a few outliers. This study introduces an efficient algorithm, called Spherically Normalized SVD, for robust SVD approximation that is highly insensitive to outliers, computationally scalable, and provides accurate approximations of singular vectors. The proposed algorithm achieves remarkable speed by utilizing only two applications of a standard reduced-rank SVD algorithm to appropriately scaled data, significantly outperforming competing algorithms in computation times. To assess the robustness of the approximated singular vectors and their subspaces against data contamination, we introduce new notions of breakdown points for matrix-valued input, including row-wise, column-wise, and block-wise breakdown points. Theoretical and empirical analyses demonstrate that our algorithm exhibits higher breakdown points compared to standard SVD and its modifications. We empirically validate the effectiveness of our approach in applications such as robust low-rank approximation and robust principal component analysis of high-dimensional microarray datasets. Overall, our study presents a highly efficient and robust solution for SVD approximation that overcomes the limitations of existing algorithms in the presence of outliers.

1 INTRODUCTION

Singular Value Decomposition (SVD) is one of the most useful tools in machine learning, used in processing image, video and natural languages, constructing recommender systems, and statistical data analysis. In particular, SVD is used for dimension reduction for downstream machine learning tasks, which often improves the overall performance of the task with reduced computational complexity. However, real-world data often contain noise, outliers, and other anomalies, and with only a contaminated data matrix at hand, standard SVD may provide undesirable low-rank decomposition. As a result, there is a need for robust SVD algorithms that can handle these challenges and provide accurate and robust results in the presence of data irregularities.

In this work, we introduce a new approach called Spherically Normalized SVD (SpSVD), which aims to handle outliers more effectively compared to the classical SVD and competing robust SVD algorithms including Zhang et al. (2013); Candès et al. (2011); Brahma et al. (2017); Rahmani and Atia (2017). Inspired by a robust PCA proposal of Locantore et al. (1999), the SpSVD algorithm adopts the spherical normalization approach of Locantore et al. to approximate both left and right singular vectors. While the normalization gives highly robust approximations of those vectors, we additionally solve a simple optimization problem for more accurate low-rank approximation. Our algorithm is easy to implement and extremely fast to compute. Specifically, it requires a computational complexity similar to that of classical SVD for low-rank approximations. We also establish that the algorithm provides a statistically accurate approximation of singular vectors, even in the presence of infinitesimal contamination of a considerable scale.

The robustness of SpSVD is carefully evaluated by extending the notion of breakdown point, a quantitative measure of robustness, commonly used in robust estimation literature such as Tyler et al. (2023); Tang and Phillips (2016); Huber (2011). The classical definition

of observation-wise breakdown point is generalized to handle singular vectors (unit vectors) and the subspaces spanned by those and also for the cases where contamination occurs for rows, columns or individual elements of input matrix \mathbf{X} . Utilizing the generalized notion of breakdown point, we show that the singular vectors approximated by SpSVD have higher breakdown points than the classical SVD and its seemingly robust variants (Gabriel and Zamir, 1979; Liu et al., 2003; Ke and Kanade, 2005; Zhang et al., 2013). To the best of our knowledge, this work is the first study investigating breakdown points specifically related to the singular vectors and their subspaces, providing novel insights into the robustness of SVD.

The accuracy, robustness, and computational times of SpSVD are empirically compared with existing robust SVD approaches including Zhang et al. (2013); Candès et al. (2011); Brahma et al. (2017); Rahmani and Atia (2017), via simulated experiments. In particular, our proposal is on par with the best-performing algorithm—the Robust PCA (RPCA) approach of Candès et al. (2011)—in terms of the accuracy and robustness, but is up to 500 times faster (in real computation times) than RPCA, rendering its effectiveness especially for large-scale data analysis.

Related Works There have been many proposals for robust SVD or PCA (Principal Component Analysis), which can be broadly categorized into four approaches. (1) Optimization for low-rank approximation: One approach to low-rank approximation is achieved by solving optimization problems (Markopoulos et al., 2014; Barrodale, 1968; Ding et al., 2019). Specifically, Candès et al. (2011) proposed an optimization problem where a data matrix is decomposed into a low-rank approximation and a sparse outlier matrix. This approach has many variants (Wright et al., 2013; Zhou et al., 2010; Xu et al., 2010). One of them, She et al. (2016); Brahma et al. (2017) proposed using the orthogonal complement of the low-rank approximation to capture outliers lying the orthogonal complement space. (2) Minimization of element-wise loss: Finding SVD can be recasted to a minimization problem with element-wise loss. To modify SVD to be more robust, several approaches including Gabriel and Zamir (1979); Liu et al. (2003); Ke and Kanade (2005); Zhang et al. (2013) have been proposed to replace the element-wise loss with other loss functions such as L_1 -loss or Huber’s loss function. (3) Outlier filtering: In a natural attempt to achieve robustness, many researchers have proposed outlier filtering methods (Xu et al., 2012; Dikonikolas et al., 2023; Rahmani and Atia, 2017; Kong et al., 2020; Jambulapati et al., 2020) including robust covariance estimation by subsampling techniques. However, most robust covariance matrix estimations

including Campbell (1980); Rousseeuw (1984, 1985); Hubert et al. (2018) require large sample size ($n \gg p$), and are not generally not applicable to SVD. (4) Projection pursuits: In that the first principal component is the direction maximizing the variability of data, researchers including Croux et al. (2013, 2007) proposed approaches based on projection pursuit where the goal is to find the direction that maximizes the dispersity of data points for some robust dispersity measure, such as the first quantile of the pairwise differences. Most of these approaches are not applicable to large-scale data analysis, due to its high computational costs. We have numerically compared our proposal with most of the approaches above, but chose not to present the results because of either poor performances or excessively long computation times.

Most of the aforementioned algorithms do not provide robust SVD, but robust PCA, which primarily concentrate on subspace recovery. Such methods may robustly recover right singular vectors, but do not provide robust left singular vectors and singular values. Note that there are data-analytic situations in which both the right and left singular vectors are simultaneously needed, for instance in multi-source data analysis (Feng et al., 2018; Lock et al., 2013; Prothero et al., 2022).

Hampel (1968) and Huber (1992) introduced the notion of breakdown point, a quantitative measure of robustness in the presence of distributional contamination. They also introduced a probabilistic measure of robustness via the influence function. Some researchers have investigated robustness of unit vectors via influence functions Ko and Guttorp (1988); Ko and Chang (1993). Our work is the first to extend the notion of breakdown points for functions taking values in the unit sphere and Grassmannian manifolds.

2 PROPOSED ALGORITHM: Spherically Normalized SVD

Background The SVD of a $n \times p$ real-valued matrix \mathbf{X} is denoted by $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^T = \sum_{r=1}^{n \wedge p} d_r \mathbf{u}_r \mathbf{v}_r^T$, where the diagonal matrix \mathbf{D} contains the non-negative singular values d_r , arranged in descending order, and \mathbf{u}_r and \mathbf{v}_r are the r th left and right singular vectors, respectively, corresponding to the r th largest singular value d_r . To motivate our construction of robust SVD algorithms, we view the real-valued matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T \in \mathbb{R}^{n \times p}$ as a data matrix, collecting the observed values of p variables from n individuals.

The principal component analysis (PCA) applied to the data matrix \mathbf{X} is closely related to the SVD of \mathbf{X} . Assuming that \mathbf{X} is column-centered, the r th (empirical) principal component (PC) direction vector is

the r th right singular vector \mathbf{v}_r . The r th PC scores are given by the n -vector $\mathbf{X}\mathbf{v}_r = (\mathbf{x}_1^T \mathbf{v}_r, \dots, \mathbf{x}_n^T \mathbf{v}_r)^T$, consisting of the projection of each data point onto \mathbf{v}_r . In fact, the PC score vector is also given by the SVD, that is, $\mathbf{X}\mathbf{v}_r = d_r \mathbf{u}_r$, in which the unit vector \mathbf{u}_r collects the standardized PC scores for n individuals, and the sample standard deviation of the PC scores is d_r/\sqrt{n} . Dimensionality reduction in data matrix \mathbf{X} can be equivalently achieved either by a rank- R SVD approximation $\tilde{\mathbf{X}}_R^{\text{svd}} := \sum_{r=1}^R d_r \mathbf{u}_r \mathbf{v}_r^T$, or by collecting the first R triples of PC direction, score and standard deviation. Note that $\tilde{\mathbf{X}}_R^{\text{svd}}$ is the best rank- R approximation to \mathbf{X} in terms of the Frobenius norm.

Motivation Both the SVD and standard PCA estimates are highly sensitive to data contamination in the data matrix \mathbf{X} . Consider an original data matrix \mathbf{X} and a contaminated version \mathbf{Z} . In the contaminated version, the values of the first row \mathbf{x}_1^T of \mathbf{X} are replaced with arbitrary values. Denoting $\mathbf{v}_r(\mathbf{X})$ and $\mathbf{v}_r(\mathbf{Z})$ as the r th right singular vectors of \mathbf{X} and \mathbf{Z} respectively, we demonstrate in Section 3 that the singular vectors are sensitive to their input. Even if there is only one differing row (data point) between \mathbf{X} and \mathbf{Z} , the difference between $\mathbf{v}_1(\mathbf{X})$ and $\mathbf{v}_1(\mathbf{Z})$ can be substantial. This occurs due to the presence of arbitrarily large $\|\mathbf{z}_1\|_2$ values, causing the singular vector $\mathbf{v}_1(\mathbf{Z})$ to align almost parallel to \mathbf{z}_1 . To limit such a potentially massive contribution of a single observation, Locantore et al. (1999) proposed to *normalize* each of n data points in \mathbf{X} , transforming \mathbf{x}_i to $\mathbf{x}_i/\|\mathbf{x}_i\|_2$, then to apply the standard PCA algorithm for the normalized data, for their proposal of a robust PCA. Since all data points are on the unit sphere in \mathbb{R}^p after normalization, the contribution of potential outliers is naturally limited. Building upon the robust PCA approach of Locantore et al. (1999), we propose a novel technique called *Spherically Normalized SVD algorithm* (SpSVD for short), which provides a robust approximation of the first R left and right singular vectors and singular values of (uncontaminated) \mathbf{X} , obtained purely from potentially contaminated matrix \mathbf{Z} .

Algorithm Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T \in \mathbb{R}^{n \times p}$ be a potentially contaminated data matrix. For a predetermined rank $R \in \{1, \dots, n \wedge p\}$ our goal is to define the ordered triple $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$, for $r = 1, \dots, R$, that provides a highly insensitive and accurate rank- R approximation $\tilde{\mathbf{X}}_R^{\text{SP}} := \sum_{r=1}^R d_r^{\text{SP}} \mathbf{u}_r^{\text{SP}} (\mathbf{v}_r^{\text{SP}})^T$ of \mathbf{X} .

For the approximation of right singular vectors, we individually scale each row of matrix \mathbf{X} to have unit length. Let $\tilde{\mathbf{X}}_{\text{row}} = [\mathbf{x}_1/\|\mathbf{x}_1\|_2, \dots, \mathbf{x}_n/\|\mathbf{x}_n\|_2]^T$ represent the row-normalized data matrix. Subsequently, a standard low-rank SVD algorithm is applied to $\tilde{\mathbf{X}}_{\text{row}}$ to obtain the R right singular vec-

tors corresponding to the R largest singular values of $\tilde{\mathbf{X}}_{\text{row}}$. The set of these right singular vectors is denoted as $V^R = \{\mathbf{v}_1(\tilde{\mathbf{X}}_{\text{row}}), \dots, \mathbf{v}_R(\tilde{\mathbf{X}}_{\text{row}})\}$. Similarly, for the approximation of left singular vectors, we scale and collect each *column* of \mathbf{X} in the column-normalized data matrix $\tilde{\mathbf{X}}^{\text{col}}$. The first R left singular vectors of $\tilde{\mathbf{X}}^{\text{col}}$ are then collected in the set $U^R = \{\mathbf{u}_1(\tilde{\mathbf{X}}^{\text{col}}), \dots, \mathbf{u}_R(\tilde{\mathbf{X}}^{\text{col}})\}$. While the elements in V^R and U^R are candidates for $(\mathbf{v}_r^{\text{SP}}$ and $\mathbf{u}_r^{\text{SP}})$, respectively, we do not set $(\mathbf{u}_r(\tilde{\mathbf{X}}^{\text{col}}), \mathbf{v}_r(\tilde{\mathbf{X}}_{\text{row}}))$ for $(\mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$. This is because using mismatched labels can provide a better approximation of \mathbf{X} .

With V^R and U^R at hand, the triple $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$ is defined sequentially. For the first triple, we solve the following:

$$(d_1^{\text{SP}}, \mathbf{u}_1^{\text{SP}}, \mathbf{v}_1^{\text{SP}}) = \underset{d \in \mathbb{R}, \mathbf{u} \in U^R, \mathbf{v} \in V^R}{\operatorname{argmin}} \|\mathbf{X} - d\mathbf{u}\mathbf{v}^T\|_{\text{F}_1}, \quad (1)$$

where $\|\mathbf{A}\|_{\text{F}_1} = \sum_{i,j} |a_{ij}|$ is the element-wise 1-norm of the matrix \mathbf{A} . For a fixed pair of $\mathbf{u} = (u_1, \dots, u_n)^T \in U^R$ and $\mathbf{v} = (v_1, \dots, v_p)^T \in V^R$, finding the solution to (1) with respect to $d \in \mathbb{R}$ is equivalent to solving a weighted median problem,¹

$$\min_d \sum_{u_i \neq 0, v_j \neq 0} |u_i v_j| \frac{x_{ij}}{u_i v_j} - d|. \quad (2)$$

For $r = 2, \dots, R$, optimization problems similar to (1) are used to define the r th triple $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$, but with the first $r-1$ triples deflated from each of \mathbf{X} , U^R and V^R . That is, we solve

$$(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}}) = \underset{d \in \mathbb{R}, \mathbf{u} \in U_r^R, \mathbf{v} \in V_r^R}{\operatorname{argmin}} \|\mathbf{X}_r - d\mathbf{u}\mathbf{v}^T\|_{\text{F}_1}, \quad (3)$$

where $\mathbf{X}_r = \mathbf{X} - \sum_{l=1}^{r-1} d_l^{\text{SP}} \mathbf{u}_l^{\text{SP}} (\mathbf{v}_l^{\text{SP}})^T$, $U_r^R = U^R \setminus \{\mathbf{u}_1^{\text{SP}}, \dots, \mathbf{u}_{r-1}^{\text{SP}}\}$, and $V_r^R = V^R \setminus \{\mathbf{v}_1^{\text{SP}}, \dots, \mathbf{v}_{r-1}^{\text{SP}}\}$. Note that in (1) and (3) above, the approximated singular value d_r^{SP} may be negative. Since singular values are, by definition, non-negative, for each $r = 1, \dots, R$, we update $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$ by $(s_r d_r^{\text{SP}}, s_r \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$, if $s_r := \operatorname{sign}(d_r^{\text{SP}}) \neq 0$. This procedure is summarized in Algorithm 1.

We remark that one may try to choose $(\mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$ to be

$$(\mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}}) := (\mathbf{u}_r(\tilde{\mathbf{X}}^{\text{col}}), \mathbf{v}_r(\tilde{\mathbf{X}}_{\text{row}})), \quad (4)$$

and solve (1) and (3) only with respect to d . Since the candidates in V^R and U^R are from different normalizations, there is no compelling reason to believe

¹The weighted median problem can be efficiently solved by the median of medians algorithm (Blum et al., 1973), or by other selection algorithms (Cormen et al., 2022). Therefore, by solving the weighted median for all candidate pairs, we obtain a solution to (1).

Algorithm 1: Rank- R Approximation by SpSVD

- Input:** data matrix
 $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T = (\mathbf{x}^1, \dots, \mathbf{x}^p) \in \mathbb{R}^{n \times p}$, $R \in \{1, \dots, n \wedge p\}$.
- Output:** $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$ for $r = 1, \dots, R$
- 1 Normalize $\tilde{\mathbf{X}}_{\text{row}} = [\mathbf{x}_1/\|\mathbf{x}_1\|_2, \dots, \mathbf{x}_n/\|\mathbf{x}_n\|_2]^T$ and $\tilde{\mathbf{X}}^{\text{col}} = [\mathbf{x}^1/\|\mathbf{x}^1\|_2, \dots, \mathbf{x}^p/\|\mathbf{x}^p\|_2]$
 - 2 Apply SVD to $\tilde{\mathbf{X}}_{\text{row}}$ to obtain rank- R approximation $\sum_{r=1}^R \tilde{d}_r \tilde{\mathbf{u}}_r \left(\mathbf{v}_r(\tilde{\mathbf{X}}_{\text{row}}) \right)^T$
 - 3 Apply SVD to $\tilde{\mathbf{X}}^{\text{col}}$ to obtain rank- R approximation $\sum_{r=1}^R \tilde{c}_r \left(\mathbf{u}_r(\tilde{\mathbf{X}}^{\text{col}}) \right) \tilde{\mathbf{v}}_r^T$
 - 4 Set $V^R = \{\mathbf{v}_1(\tilde{\mathbf{X}}_{\text{row}}), \dots, \mathbf{v}_R(\tilde{\mathbf{X}}_{\text{row}})\}$ and $U^R = \{\mathbf{u}_1(\tilde{\mathbf{X}}^{\text{col}}), \dots, \mathbf{u}_R(\tilde{\mathbf{X}}^{\text{col}})\}$
 - 5 **for** $r = 1, \dots, R$ **do**
 - 6 Find the r th solution $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$ of (3)
 - 7 Update $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$ by $(s_r d_r^{\text{SP}}, s_r \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$, if $s_r := \text{sign}(d_r^{\text{SP}}) \neq 0$
-

that the orders in the left and right singular values are related as in (4). We have empirically found that the solutions of (1) and (3) are typically different from the naive choice (4), more often when R is large. Our approach of searching over the R^2 pairs of candidates provides generally better rank- R approximations than the naive choice. On the other hand, potential downsides of our approach include the following: The r th triple $(d_r^{\text{SP}}, \mathbf{u}_r^{\text{SP}}, \mathbf{v}_r^{\text{SP}})$ may depend on the choice of the rank $R \geq r$, and the computational complexity increases because of the discrete optimization needed in (3). Nevertheless, for small R the overall computational complexity of the proposed SpSVD algorithm is similar to that of a standard low-rank SVD algorithm.

Computational Complexity The computational complexity of an algorithm is a measure of the amount of resources required for solving a problem of a given size. The overall computational complexity of Algorithm 1 applied to rank- R approximation of a real matrix of size $n \times p$ is $O(npR^3)$. While the two applications of SVD as well as the normalization require $O(npR)$, the algorithm solves the weighted median problem (2) in the linear time complexity of $O(np)$ for $R(R+1)(2R+1)/3$ times. As a comparison, the computational complexity of the standard rank- R SVD algorithm is $O(npR)$ (Xu et al., 2023; Yi et al., 2016; Shamir, 2016; Allen-Zhu and Li, 2016). Algorithm 1 is as efficient as SVD when R is small. As we compare numerically in Section 5, computing the approximation of a low-rank SVD by our approach is up to 500 times (on average in real computation times) faster

than state-of-the-art robust SVD algorithms proposed in Zhang et al. (2013); Brahma et al. (2017); Candès et al. (2011).

Statistical Accuracy When the data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ is viewed as a collection of n observations \mathbf{x}_i , the right singular vector $\mathbf{v}_r(\mathbf{X})$ is equivalent to the eigenvector of the sample covariance matrix of \mathbf{X} . Treating \mathbf{v}_r^{SP} as an estimator of the r th eigenvector of the population covariance matrix, the estimator \mathbf{v}_r^{SP} is consistent under adequate assumptions; see Section 7.1.1 of the supplementary material. While the standard SVD also provides a consistent estimator, SpSVD exhibit a statistical accuracy over contamination. Let \mathcal{F}_{Σ} denote a mean-zero, p -dimensional elliptical distribution (Cambanis et al., 1981), with covariance matrix Σ , whose eigen-decomposition is given by $\Sigma = \sum_{j=1}^p \lambda_j \mathbf{v}_j \mathbf{v}_j^T$. Assume that $\lambda_{j+1} > \lambda_j$ for all j . Let $\epsilon > 0$ be the fraction of contaminations among n samples: That is, the samples consist of both i.i.d (uncontaminated) samples $\mathbf{x}_1, \dots, \mathbf{x}_{(1-\epsilon)n} \sim \mathcal{F}_{\Sigma}$ and the outliers $\mathbf{y}_1, \dots, \mathbf{y}_{\epsilon n}$ of arbitrary size and directions. Let \mathbf{v}_j^{SP} be the j th left singular vector obtained by the proposed SpSVD applied to $[\mathbf{x}_1, \dots, \mathbf{x}_{(1-\epsilon)n}, \mathbf{y}_1, \dots, \mathbf{y}_{\epsilon n}]^T$.

Theorem 1 (Statistical accuracy over infinitesimal contamination). *If (i) $n \geq Cp/\epsilon$, for a constant $C > 0$, and (ii) for $j = 1, \dots, p$, $\delta_j := \{|d_{j+1} - d_j|, |d_j - d_{j-1}|\} > 0$ where d_j is the j th singular value of the covariance matrix of $\frac{\mathbf{x}_1}{\|\mathbf{x}_1\|_2}$, then*

$$\begin{aligned} \mathbb{E}[\min\{\|\mathbf{v}_j^{\text{SP}} - \mathbf{v}_j\|_2, \|\mathbf{v}_j^{\text{SP}} + \mathbf{v}_j\|_2\}] \\ \leq \frac{1}{\delta_j} (C'\epsilon + C''\sqrt{(1-\epsilon)\epsilon}) \end{aligned} \quad (5)$$

for some absolute constants $C', C'' > 0$.

In the above theorem, as the fractions ϵ and $\frac{p}{n}$ go to zero, the singular vectors of our method converge to the target singular vectors in an appropriate statistical context. The assumption (i) $n \geq Cp/\epsilon$ is more relaxed than the assumption “ $n \geq Cp/\epsilon^2$ ” provided in Diakonikolas et al. (2023) (Similar assumptions were imposed in Jambulapati et al. (2020); Kong et al. (2020).) If we assume $n \geq Cp/\epsilon^2$, the second term of (5) becomes negligible. Our result states the accuracy of unit singular vectors \mathbf{v}_j^{SP} directly compared to \mathbf{v}_j (for all $j = 1, \dots, p$). This is in contrast to the statements in Xu et al. (2012); Jambulapati et al. (2020); Diakonikolas et al. (2023), in which the accuracy of the only the first vector $\hat{\mathbf{v}}_1$ is compared indirectly by bounding $|\hat{\mathbf{v}}_1^T \Sigma \hat{\mathbf{v}}_1 - \lambda_1|$.

3 EXTENSIONS OF BREAKDOWN POINTS

The breakdown point, originally proposed by Hampel (1968), and studied by Huber and Donoho (1983); Huber (1984, 2011), is a common tool for evaluating quantitative robustness of statistics. Viewing a real-valued statistic as a function $f : \mathcal{X}^n \rightarrow \mathbb{R}$ that takes as input n data points $\mathbf{X} := (\mathbf{x}_1, \dots, \mathbf{x}_n) \in \mathcal{X}^n$ and outputs a real-valued $f(\mathbf{X})$, the breakdown point of f at the given data \mathbf{X} is defined as the minimum number of corrupted data points that cause the statistic to “break down.” Formally, the breakdown point of f at \mathbf{X} is

$$\text{bp}(f; \mathbf{X}) := \min_{1 \leq l \leq n} \{l : \sup_{\mathbf{Z}_l} |f(\mathbf{Z}_l) - f(\mathbf{X})| = \infty\}, \quad (6)$$

where the supremum is taken over all possible corrupted collections \mathbf{Z}_l that are obtained from \mathbf{X} by replacing l data points of \mathbf{X} with arbitrary values. In this sense, a function f (giving the value of a statistic $f(\mathbf{X})$) is said to *break down* if the difference between the statistics computed from corrupted data and from the original data, i.e., $f(\mathbf{Z}_l)$ and $f(\mathbf{X})$, is as large as possible, which was defined to be the infinity in the original definition (6) of the breakdown point. Note that in the literature (Huber and Donoho, 1983; Huber, 1984, 2011; Lopuhaa and Rousseeuw, 1991), the (finite-sample) breakdown point is in fact defined as $\text{bp}(f; \mathbf{X})/n$, the fraction of the number of corrupted data points and the sample size. Nevertheless, for notational simplicity, we regard the number of data points n as fixed, and define the breakdown point as a whole number. The breakdown point $\text{bp}(f; \mathbf{X})$ represents a critical threshold where a breakdown of f does not occur when the number of corrupted data points is below $\text{bp}(f; \mathbf{X})$, but breakdown can occur when the number of corrupted data points is equal to or exceeds $\text{bp}(f; \mathbf{X})$.

We extend the notion of breakdown points to the situations where the singular vectors $\mathbf{u}_r(\mathbf{X})$ and $\mathbf{v}_r(\mathbf{X})$ and the subspaces spanned by these are the statistics of interest. This involves two distinct considerations: Breakdown of unit-sphere S^{k-1} and grassmannian $\text{Gr}(k, r)$ -valued statistics, and breakdown with respect to contamination matrices.

Breakdown of Unit Vectors and Subspaces

Viewing each of the singular vectors $\mathbf{u}_r(\mathbf{X})$ and $\mathbf{v}_r(\mathbf{X})$ (or $\mathbf{u}_r^{\text{SP}}(\mathbf{X})$ and $\mathbf{v}_r^{\text{SP}}(\mathbf{X})$) as a statistic, the notion of breakdown point (6) naturally applies with the following modifications. Since a singular vector lies in the unit sphere S^{k-1} (for $k = n$ or p), we measure the difference between $\mathbf{v}, \mathbf{w} \in S^{k-1}$ by $\theta(\mathbf{v}, \mathbf{w}) := \arccos(|\mathbf{v}^T \mathbf{w}|)$, the “angle” between two directions. The maximum difference in this case is $\pi/2$. We will also be interested in the

subspaces spanned by singular vectors. Let $\mathcal{V} = \mathcal{V}(\mathbf{X})$ be the subspace spanned by $\{\mathbf{v}_r(\mathbf{X}) : r = 1, \dots, R\}$. Then $\mathcal{V} \in \text{Gr}(k, R)$, the Grassmannian manifold consisting of R -dimensional subspaces in \mathbb{R}^k . For two subspaces $\mathcal{V} = \text{span}(\mathbf{V}), \mathcal{W} = \text{span}(\mathbf{W}) \in \text{Gr}(k, R)$, the difference may be measured via the largest canonical angle $\theta(\mathcal{V}, \mathcal{W}) := \arccos(d_{\min}(\mathbf{V}^T \mathbf{W}))$, where $d_{\min}(\mathbf{A})$ is the smallest singular value of \mathbf{A} . Note that for $R = 1$, $\theta(\text{span}(\mathbf{v}), \text{span}(\mathbf{w})) = \theta(\mathbf{v}, \mathbf{w})$. The maximum difference is also $\pi/2$, and we say $\mathcal{V} : \mathbb{R}^{n \times p} \rightarrow \text{Gr}(k, R)$ (or S^k) breaks down at \mathbf{X} by replacing l data points, if $\sup_{\mathbf{Z}_l} \theta(\mathcal{V}(\mathbf{Z}_l), \mathcal{V}(\mathbf{X})) = \pi/2$. The breakdown point of \mathcal{V} at \mathbf{X} is then given by (6) with the definition of “breakdown” given above.

Breakdown with respect to Contamination of Matrices

A data matrix, to which SVD is performed, is not always a statistical data matrix in $\mathbb{R}^{n \times p}$ consisting of n data points with p variables. We generalize the mechanism of data contamination from the observation-wise data contamination (6) to three different types of data contaminations. Treating each row, column, or element as a data point, we will discuss row-wise, column-wise, and block-wise contaminations of data matrix \mathbf{X} .

Let $\mathcal{V} : \mathbb{R}^{n \times p} \rightarrow \text{Gr}(k, r)$ be a statistic of interest. The row-wise breakdown point measures the robustness of \mathcal{V} in terms of contamination of the rows of input matrix \mathbf{X} , and coincides with the traditional notion of breakdown point (6) when rows represent data points. That is, we define $\text{bp}_{\text{row}}(\mathcal{V}; \mathbf{X}) = \min\{l : \sup_{\mathbf{Z}_l} \theta(\mathcal{V}(\mathbf{Z}_l), \mathcal{V}(\mathbf{X})) = \pi/2, 1 \leq l \leq n\}$, where the supremum is taken over all possible \mathbf{Z}_l obtained by replacing l rows of \mathbf{X} by arbitrary values. Similarly, the column-wise breakdown point measures the robustness of \mathcal{V} with respect to contaminated columns (viewing each column as a data point), and is $\text{bp}_{\text{col}}(\mathcal{V}; \mathbf{X}) = \min\{l : \sup_{\mathbf{Z}^l} \theta(\mathcal{V}(\mathbf{Z}^l), \mathcal{V}(\mathbf{X})) = \pi/2, 1 \leq l \leq p\}$, in which \mathbf{Z}^l is given by replacing l columns of \mathbf{X} .

In situations where each element of the $n \times p$ matrix \mathbf{X} is considered as an observation, outlying observations may be scattered across the matrix, and it becomes challenging to devise an informative notion of robustness. We focus on the case that contamination occurs within a (possibly non-consecutive) block. For example, if three elements $x_{1,1}$, $x_{1,3}$ and $x_{2,3}$ of \mathbf{X} are contaminated, then we say the outliers lie in a block of size $(2, 2)$, in which the numbers correspond to two rows and two columns, respectively. We say \mathcal{V} breaks down at block-size (k, l) (at \mathbf{X}) if $\sup_{\mathbf{Z}_k^l} \theta(\mathcal{V}(\mathbf{Z}_k^l), \mathcal{V}(\mathbf{X})) = \pi/2$, where the corrupted data \mathbf{Z}_k^l are given by replacing the elements in a $k \times l$ block of \mathbf{X} . Recall that the breakdown point $\text{bp}(f; \mathbf{X})$ (6) is a threshold, i.e., the minimum number of data points needed for the statistic

to break down. To extend the definition of breakdown point to this block-wise contamination scenario, we adopt a partial order relation “ \prec ” among the block-sizes in $\mathcal{B}_{n,p} = \{(k, l) : 1 \leq k \leq n, 1 \leq l \leq p\}$, given by (i) $(i, j) \preceq (k, l)$ if $i \leq k$ and $j \leq l$, and (ii) $(i, j) \prec (k, l)$ if $(i, j) \preceq (k, l)$ and $(i, j) \neq (k, l)$.

The set $(\mathcal{B}_{n,p}, \prec)$ is only partially ordered, meaning that there are block-sizes (i, j) and (k, l) that can not be ordered; take $(2, 3)$ and $(3, 1)$ as an example. This is unavoidable due to the two-dimensional nature of $\mathcal{B}_{n,p}$. Nevertheless, utilizing the partial order provides a definition for block-wise breakdown point as a “tipping” point.

Definition 1. *We say that \mathcal{V} has a block-wise breakdown point (i, j) at \mathbf{X} if (i) for any $(k, l) \succeq (i, j)$, \mathcal{V} breaks down at block-size (k, l) , and (ii) for any $(k', l') \prec (i, j)$, \mathcal{V} does not break down at block-size (k', l') .*

It is possible that there are multiple block-wise breakdown points for \mathcal{V} , and we denote the set of all block-wise breakdown points of \mathcal{V} at \mathbf{X} by $\text{BP}(\mathcal{V}; \mathbf{X})$. The notion of block-wise breakdown points is more powerful than row and column-wise breakdown points, as the following lemma states.

Lemma 2. *(i) If $(k, 1) \in \text{BP}(\mathcal{V}; \mathbf{X})$ for some $1 \leq k \leq n$, then $\text{bp}_{\text{col}}(\mathcal{V}; \mathbf{X}) = 1$, and (ii) if $(1, l) \in \text{BP}(\mathcal{V}; \mathbf{X})$ for some $1 \leq l \leq p$, then $\text{bp}_{\text{row}}(\mathcal{V}; \mathbf{X}) = 1$.*

4 ROBUSTNESS OF SpSVD

We begin by highlighting the lack of robustness of the standard SVD. Let $\mathcal{V}_R : \mathbb{R}^{n \times p} \rightarrow \text{Gr}(p, R)$ be the function that gives the rank- R right singular subspace $\mathcal{V}_R(\mathbf{X}) = \text{span}(\mathbf{v}_1(\mathbf{X}), \dots, \mathbf{v}_R(\mathbf{X}))$ of input matrix \mathbf{X} . The function \mathcal{U}_R for the left singular subspace is similarly defined.

Proposition 3. *Let \mathbf{X} be any $n \times p$ real matrix. The following holds for any $R = 1, \dots, n \wedge p$.*

- (i) *For some $k \leq R + 1$, $(k, 1) \in \text{BP}(\mathcal{U}_R; \mathbf{X})$. In particular, $\text{bp}_{\text{row}}(\mathcal{U}_R; \mathbf{X}) \leq R + 1$, and $\text{bp}_{\text{col}}(\mathcal{U}_R; \mathbf{X}) = 1$.*
- (ii) *For some $l \leq R + 1$, $(1, l) \in \text{BP}(\mathcal{V}_R; \mathbf{X})$. In particular, $\text{bp}_{\text{row}}(\mathcal{V}_R; \mathbf{X}) = 1$, and $\text{bp}_{\text{col}}(\mathcal{V}_R; \mathbf{X}) \leq R + 1$.*

In particular, the first right singular vector $\mathbf{v}_1(\mathbf{X})$ breaks down even with contamination of one row of \mathbf{X} , or two elements in a block of size $(2, 1)$. This shows that the standard SVD is highly sensitive to outliers in the matrix. Surprisingly, a family of seemingly robust algorithms for SVD approximation, which we call Element-wise Loss SVD (or ELSVD for short), turns out to have very low breakdown points. In particular, the ELSVD, studied in Gabriel and Zamir (1979);

Liu et al. (2003); Ke and Kanade (2005); Zhang et al. (2013), is given by solving the following problem: For $r = 1, \dots, R$,

$$(d_r^\rho, \mathbf{u}_r^\rho, \mathbf{v}_r^\rho) = \underset{d \geq 0, \mathbf{u} \in S^{n-1}, \mathbf{v} \in S^{p-1}}{\text{argmin}} \rho(\mathbf{X} - d\mathbf{u}\mathbf{v}^T) + \mathcal{P}_1(\mathbf{u}) + \mathcal{P}_2(\mathbf{v}), \quad (7)$$

subject to an orthogonality constraint. In (7), the function $\rho(\cdot)$ is defined as $\rho(\mathbf{Z}) = \sum_{i=1}^p \sum_{j=1}^n \rho_{ij}(z_{ij})$ for $\mathbf{Z} = (z_{ij})$, where $\rho_{ij} : \mathbb{R} \rightarrow \mathbb{R}$ are symmetric and non-negative loss functions, and $\mathcal{P}_1, \mathcal{P}_2$ are regularization terms. Note that for $\rho_{ij}(z) = z^2$ without the regularization terms, the solution to (7) coincides with the standard SVD. The element-wise losses can be set as the L_1 or Huber loss, which are commonly believed to induce robustness. Let $\mathcal{U}_R^\rho : \mathbb{R}^{n \times p} \rightarrow \text{Gr}(n, R)$ and $\mathcal{V}_R^\rho : \mathbb{R}^{n \times p} \rightarrow \text{Gr}(p, R)$ be the functions that provide the R th left and right singular spaces, respectively, obtained from ELSVD using some ρ, \mathcal{P}_1 , and \mathcal{P}_2 .

Theorem 4. *Suppose that ρ_{ij} satisfies that $\rho_{ij}(z) \rightarrow \infty$ as $|z| \rightarrow \infty$ (for all i, j), and the functions $\mathcal{P}_1, \mathcal{P}_2$ restricted to the domains $\mathbf{u} \in S^{n-1}$ and $\mathbf{v} \in S^{p-1}$ respectively are each upper bounded. Then, the conclusions of Proposition 3 hold when \mathcal{U}_R and \mathcal{V}_R are each replaced by \mathcal{U}_R^ρ and \mathcal{V}_R^ρ .*

In contrast, the SVD approximations given by the proposed SpSVD have higher breakdown points, as we explain below. Let $\mathcal{V}_R^{\text{Sp}} : \mathbb{R}^{n \times p} \rightarrow \text{Gr}(p, R)$ be given by $\mathcal{V}_R^{\text{Sp}}(\mathbf{X}) = \text{span}(\mathbf{v}_1^{\text{Sp}}, \dots, \mathbf{v}_R^{\text{Sp}})$, the dimension- R right singular subspace approximated by SpSVD. The left singular subspace function $\mathcal{U}_R^{\text{Sp}}$ is defined similarly.

Recall that $\tilde{\mathbf{X}}_{\text{row}} = (\mathbf{x}_1/\|\mathbf{x}_1\|_2, \dots, \mathbf{x}_n/\|\mathbf{x}_n\|_2)^T \in \mathbb{R}^{n \times p}$ denotes the row-normalized matrix. Let $\mathbf{\Pi}_R := \mathbf{\Pi}_R(\mathbf{X})$ be the $p \times p$ matrix of projection onto the orthogonal complement of $\mathcal{V}_R^{\text{Sp}}(\mathbf{X})$, and let $\lambda_r(\mathbf{X})$ denotes the r th largest singular value of \mathbf{X} . Define

$$n_R := n_R(\mathbf{X}) = \min_{1 \leq k \leq n} \left\{ k : k \geq \inf_{\tilde{\mathbf{X}}_k} \left\{ (\lambda_R(\tilde{\mathbf{X}}_k))^2 - (\lambda_1(\tilde{\mathbf{X}}_k \mathbf{\Pi}_R))^2 \right\} \right\}, \quad (8)$$

where infimum is taken over all possible submatrices $\tilde{\mathbf{X}}_k \in \mathbb{R}^{(n-k) \times p}$ of $\tilde{\mathbf{X}}_{\text{row}}$ obtained by choosing $n - k$ rows of $\tilde{\mathbf{X}}_{\text{row}}$. Similarly, let $p_R := p_R(\mathbf{X})$ be given by $p_R(\mathbf{X}) = n_R(\mathbf{X}^T)$. The number n_R represents a lower bound on the minimum number of rows that can break down $\mathcal{V}_R^{\text{Sp}}$, since it can be shown that the inequality in (8) is a necessary condition for $\mathcal{V}_R^{\text{Sp}}$ to break down with k outliers. Similarly, the number p_R is a lower bound for the column-wise breakdown point for the left singular subspace approximation $\mathcal{U}_R^{\text{Sp}}$.

Theorem 5. *Let \mathbf{X} be any $n \times p$ real matrix. The following holds for any $R = 1, \dots, n \wedge p$.*

- (i) $\text{bp}_{\text{row}}(\mathcal{U}_R^{\text{Sp}}; \mathbf{X}) \leq R + 1$, $\text{bp}_{\text{col}}(\mathcal{U}_R^{\text{Sp}}; \mathbf{X}) \geq p_R$, and for any $(k, l) \in \text{BP}(\mathcal{U}_R^{\text{Sp}}; \mathbf{X})$, $(1, p_R) \preceq (k, l)$.
- (ii) $\text{bp}_{\text{row}}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X}) \geq n_R$, $\text{bp}_{\text{col}}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X}) \leq R + 1$, and for any $(k, l) \in \text{BP}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X})$, $(n_R, 1) \preceq (k, l)$

In Theorem 5, $(n_R, 1)$ represents a lower bound of $\text{BP}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X})$. This bound may be perceived as too low because it only includes the column size 1. However, the theorem also implies that no breakdown occurs for $\mathcal{V}_R^{\text{Sp}}$ by contamination of blocks of size (k, l) , for any $k < n_R$ and $l = 1, \dots, p$. This is because every point in $\text{BP}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X})$ is greater than or equal to $(n_R, 1)$. In other words, if $(k, k) \in \text{BP}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X})$, then $k \geq n_R$.

Evaluating the lower bounds n_R and p_R appears to be challenging. Through numerical experiments, we have observed that while n_R and p_R depend on the matrix \mathbf{X} , they tend to be larger when there is a larger gap between $\lambda_R(\mathbf{X})$ and $\lambda_{R+1}(\mathbf{X})$; See Section 7.3.5 of the supplementary material.

We additionally investigate the breakdown points in COP (Rahmani and Atia, 2017), which adopt screening out potential outliers; See Section 7.4 of the supplementary material.

5 NUMERICAL STUDIES

In this section we evaluate the empirical performance of SpSVD in terms of accuracy, robustness and computational scalability, making comparisons to the standard SVD algorithm, ELSVD with Huber’s loss of Zhang et al. (2013), RPCA of Candès et al. (2011), R2PCP of Brahma et al. (2017), as well as COP of Rahmani and Atia (2017).

RPCA aims to recover a low-rank matrix and a sparse outlier matrix from the data matrix by decomposing it into the sum of the two matrices, while R2PCP (She et al., 2016; Brahma et al., 2017) models not only a sparse outlier matrix but also the orthogonal complement of the low-rank matrix where outliers lying. Since both approaches provide robust low-rank approximations, we apply the standard SVD to the low-rank approximation to extract the left and right singular vectors and singular value approximations. COP adopts normalizing and filtering out outliers in constructing PC directions.

Simulation Experiment We model the data matrix without outliers $\mathbf{X} \in \mathbb{R}^{n \times p}$ as the sum of low-rank \mathbf{L} and \mathbf{E} consisting of standard normal random noises with $n = 200$ and $p = 100$. The low-rank \mathbf{L} is $\mathbf{L} = \sum_{r=1}^3 d_r \mathbf{u}_r \mathbf{v}_r^T$, $(d_1, d_2, d_3) = (80, 70, 60)$, and $\mathbf{U} = (\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3)$ and $\mathbf{V} = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ are randomly

sampled where the uniform distribution on the set of orthogonal matrices, respectively. We add a sparse outlier matrix \mathbf{S} with $\|\mathbf{S}\|_F = 1$ multiplied by η , a scaling parameter, to the data matrix \mathbf{X} , i.e.,

$$\mathbf{X} = \mathbf{L} + \mathbf{E}, \quad \mathbf{X}^\eta = \mathbf{L} + \eta \mathbf{S} + \mathbf{E}.$$

The outlier matrix \mathbf{S} has non-zero elements only in arbitrary blocks with the block-size $(0.05n, 0.05p) \in \mathcal{B}_{n,p}$ where the block-size is set based on the computation of the lower bounds n_R and p_R described in Theorem 5 for a reduced-size matrix; see Section 7.3.5 of the supplementary material. We evaluate the performance of different methods by gradually increasing η from 0 to 1000, and repeating the simulation 100 times for each value of η .

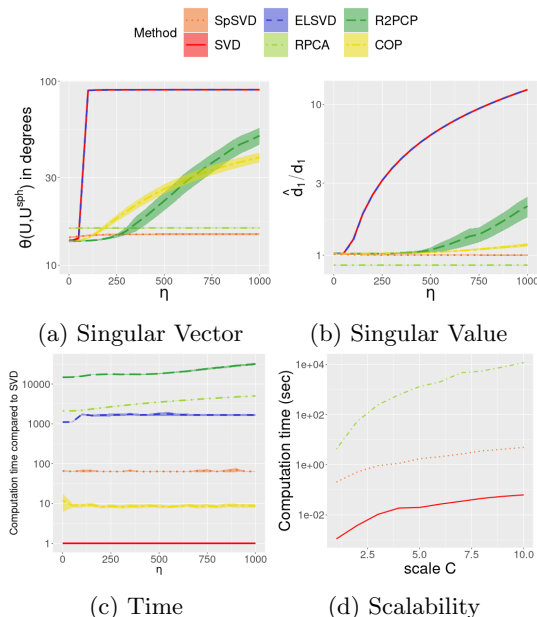


Figure 1: The Approximation Accuracy, Robustness, and Computation Times Against Increasing Magnitude η of Outliers In (a), (b), And (c). Panel (d) shows the computation times over varying size of input matrix. The y-axes are in log-scale.

Figure 1(a) illustrates the angle between 3-dimensional left singular subspace of \mathbf{L} and the approximated subspaces given by the six methods. We do not report the results for the right singular subspaces. Except that COP recovers well the right singular subspaces, the result is similar to that of the left singular subspace. Panel (b) shows the ratio of the approximated largest singular value of \mathbf{X}^η to the largest singular value of \mathbf{L} . In terms of these values, we find that SVD,

ELSVD, R2PCP, and COP do not exhibit robustness. While both SpSVD and RPCA appear to show desirable performances (with approximation errors for the singular subspace around 15° across varying η), SpSVD performs consistently better for singular value approximation ($\hat{d}_1^{\text{Sp}}/d_1 \approx 1$ and $\hat{d}_1^{\text{RPCA}}/d_1 \approx 0.86$). On the other hand, SpSVD boasts remarkably faster computation times compared to robust SVD approaches; see panel (c). Specifically, R2PCP requires approximately 500 times longer computation times on average for $\eta = 1000$.

To measure computational scalability of these algorithms, we increase the size of the matrix from 200×100 to 2000×1000 , as shown in Figure 1(d). For this comparison, we have only compared SpSVD with the standard SVD as a baseline method and RPCA, the best-performing method among the existing methods. While the computation time for SpSVD increases at a rate similar to that of SVD, the computation time for RPCA increases much faster. Overall, we find that SpSVD not only accurately approximates the SVD in the presence of massive outliers, but also is fast, requiring only 70 times longer computation times than the standard SVD, in contrast to the other methods which require more than 1000 times longer times for large-scale data. We have also experimented with a higher rank case; see Section 7.3.3 of the supplementary material.

Experiment on Gene Expression Data Matrix

Using the data consisting of gene expression levels of $n = 168$ patients with small invasive ductal carcinomas, obtained from a comparative genomic hybridization array (Gravier et al., 2010), we consider a scenario in which a block of size $(16, 16)$ is contaminated (perhaps by a physical contamination of the array). See Section 7.3.4 of the supplementary material for a detailed description of data pre-processing, and the mechanism of contamination.

The rank-2 approximation \mathbf{X}_2 of the original, uncontaminated data \mathbf{X} (via the standard SVD) is considered as a ground truth. We also obtain three rank-2 approximations of the contaminated data: by the standard SVD, denoted by $\hat{\mathbf{X}}_2^{\text{svd}}$ using the standard SVD, $\hat{\mathbf{X}}_2^{\text{Sp}}$ and $\hat{\mathbf{X}}_2^{\text{rpca}}$ obtained by SpSVD and RPCA, respectively. As expected, the SVD approximation is heavily affected by the outliers. In contrast, $\hat{\mathbf{X}}_2^{\text{Sp}}$ and $\hat{\mathbf{X}}_2^{\text{rpca}}$ provide accurate and highly robust approximations of \mathbf{X}_2 . These are graphically depicted in Figure 3 in Section 7.3.4 of the supplementary material.

To further compare the performances of SpSVD and RPCA, we repeat the above experiment for 100 times, each with different realizations of random contamination. The quality of approximation $\hat{\mathbf{X}}_2^{\text{Sp}}$ is measured by

the relative error, $r(\hat{\mathbf{X}}_2^{\text{Sp}}) := \|\mathbf{X} - \hat{\mathbf{X}}_2^{\text{Sp}}\|_{\text{F}} / \|\mathbf{X} - \mathbf{X}_2\|_{\text{F}}$. The errors $r(\hat{\mathbf{X}}_2^{\text{Sp}})$ and $r(\hat{\mathbf{X}}_2^{\text{rpca}})$ are found to be similar, with values of 1.02 and 1.01, on average, respectively, and are much smaller than $r(\hat{\mathbf{X}}_2^{\text{svd}}) \approx 51.56$. On average, SpSVD takes only 0.13 seconds for rank-2 approximation, while RPCA takes 68.56 seconds for the same task, which is more than 500 times larger than SpSVD’s computation times.

6 CONCLUSIONS AND FUTURE WORKS

In this paper, we have presented the SpSVD algorithm that provides a highly scalable, accurate, and robust approximation of SVD. To demonstrate robustness, we have extended the classical notion of breakdown point to incorporate breakdown of unit vectors and subspaces, with respect to row-wise, column-wise, and block-wise contamination of a data matrix. Using the novel notion of block-wise breakdown points, our theoretical analysis further validates that our approach produces robust singular vector approximations in the presence of block-wise contamination of input matrix contamination, outperforming the classical SVD and ELSVD. Through numerical studies, we have demonstrated not only the desirable accuracy and robustness of our approach, but also a much higher computation efficiency compared to existing methods.

We point out that our theoretical analysis can be further improved. In particular, in Theorem 5 we have only provided lower bounds of breakdown points, and these bounds are by no means optimal. Moreover, evaluating the lower bounds for a given \mathbf{X} seems very challenging, as computing n_R involves $\sum_{k=1}^{n_R} \frac{n!}{(n-k)!k!}$ comparisons, which becomes impractical for large-scale data with large values of n . A potential future direction of research is to explore alternative approaches that provide a tightened lower bound with improved computational efficiency.

Additionally, we have not theoretically examined the block-wise breakdown points for many other robust SVD methods, including RPCA of Candès et al. (2011), and R2PCP of Brahma et al. (2017). The breakdown of singular subspaces, described in terms of the maximal deviation from uncontaminated subspace, may be explored for other robust SVD methods, but appears to be technically challenging. Finally, we have not addressed the issue of rank selection, which becomes even harder with contaminated data. We leave this as a future research topic.

Acknowledgements

This work was supported by Samsung Science and Technology Foundation under Project Number SSTF-BA2002-03.

References

- Allen-Zhu, Z. and Li, Y. (2016), “LazySVD: Even faster SVD decomposition yet without agonizing pain,” *Advances in Neural Information Processing Systems*, 29.
- Barrodale, I. (1968), “ L_1 approximation and the analysis of data,” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 17, 51–57.
- Blum, M., Floyd, R. W., Pratt, V. R., Rivest, R. L., Tarjan, R. E., et al. (1973), “Time bounds for selection,” *Journal of Computer and System Sciences*, 7, 448–461.
- Brahma, P. P., She, Y., Li, S., Li, J., and Wu, D. (2017), “Reinforced robust principal component pursuit,” *IEEE Transactions on Neural Networks and Learning Systems*, 29, 1525–1538.
- Cambanis, S., Huang, S., and Simons, G. (1981), “On the theory of elliptically contoured distributions,” *Journal of Multivariate Analysis*, 11, 368–385.
- Campbell, N. A. (1980), “Robust procedures in multivariate analysis I: Robust covariance estimation,” *Journal of the Royal Statistical Society Series C: Applied Statistics*, 29, 231–237.
- Candès, E. J., Li, X., Ma, Y., and Wright, J. (2011), “Robust principal component analysis?” *Journal of the ACM (JACM)*, 58, 1–37.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2022), *Introduction to algorithms*, MIT press.
- Croux, C., Filzmoser, P., and Fritz, H. (2013), “Robust sparse principal component analysis,” *Technometrics*, 55, 202–214.
- Croux, C., Filzmoser, P., and Oliveira, M. R. (2007), “Algorithms for projection–pursuit robust principal component analysis,” *Chemometrics and Intelligent Laboratory Systems*, 87, 218–225.
- Diakonikolas, I., Kane, D., Pensia, A., and Pittas, T. (2023), “Nearly-linear time and streaming algorithms for outlier-robust PCA,” in *International Conference on Machine Learning*, PMLR, pp. 7886–7921.
- Ding, T., Zhu, Z., Ding, T., Yang, Y., Robinson, D. P., Tsakiris, M. C., and Vidal, R. (2019), “Noisy Dual Principal Component Pursuit.” in *ICML*, pp. 1617–1625.
- Feng, Q., Jiang, M., Hannig, J., and Marron, J. (2018), “Angle-based joint and individual variation explained,” *Journal of multivariate analysis*, 166, 241–265.
- Gabriel, K. R. and Zamir, S. (1979), “Lower rank approximation of matrices by least squares with any choice of weights,” *Technometrics*, 21, 489–498.
- Graciela and Fraiman (1999), “Discussion to robust principal component analysis for functional data,” *Test*, 8, 28–35.
- Gravier, E., Pierron, G., Vincent-Salomon, A., Gruel, N., Raynal, V., Savignoni, A., De Rycke, Y., Pierga, J.-Y., Lucchesi, C., Reyal, F., et al. (2010), “A prognostic DNA signature for T1T2 node-negative breast cancer patients.” *Genes, Chromosomes and Cancer*, 49, 1125–1134.
- Hampel, F. R. (1968), *Contributions to the theory of robust estimation*, University of California, Berkeley.
- Huber, P. J. (1984), “Finite sample breakdown of M - and P -Estimators,” *The Annals of Statistics*, 12, 119–126.
- (1992), “Robust estimation of a location parameter,” *Breakthroughs in Statistics: Methodology and distribution*, 492–518.
- (2011), “Robust statistics,” in *International Encyclopedia of Statistical Science*, Springer, pp. 1248–1251.
- Huber, P. J. and Donoho, D. (1983), “The notion of breakdown point,” in *A festschrift for Erich L. Lehmann*, CRC Press, pp. 157–184.
- Hubert, M., Debruyne, M., and Rousseeuw, P. J. (2018), “Minimum covariance determinant and extensions,” *Wiley Interdisciplinary Reviews: Computational Statistics*, 10, e1421.
- Jambulapati, A., Li, J., and Tian, K. (2020), “Robust sub-gaussian principal component analysis and width-independent Schatten packing,” *Advances in Neural Information Processing Systems*, 33, 15689–15701.
- Ke, Q. and Kanade, T. (2005), “Robust L_1 norm factorization in the presence of outliers and missing data by alternative convex programming,” in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, IEEE, vol. 1, pp. 739–746.
- Ko, D. and Chang, T. (1993), “Robust M-estimators on spheres,” *Journal of Multivariate Analysis*, 45, 104–136.
- Ko, D. and Guttorp, P. (1988), “Robustness of estimators for directional data,” *The Annals of Statistics*, 609–618.
- Kong, W., Somani, R., Kakade, S., and Oh, S. (2020), “Robust meta-learning for mixed linear regression with small batches,” *Advances in neural information processing systems*, 33, 4683–4696.

- Liu, L., Hawkins, D. M., Ghosh, S., and Young, S. S. (2003), “Robust singular value decomposition analysis of microarray data,” *Proceedings of the National Academy of Sciences*, 100, 13167–13172.
- Locantore, N., Marron, J., Simpson, D., Tripoli, N., Zhang, J., Cohen, K., Boente, G., Fraiman, R., Brumback, B., Croux, C., et al. (1999), “Robust principal component analysis for functional data,” *Test*, 8, 1–73.
- Lock, E. F., Hoadley, K. A., Marron, J. S., and Nobel, A. B. (2013), “Joint and individual variation explained (JIVE) for integrated analysis of multiple data types,” *The annals of applied statistics*, 7, 523.
- Lopuhaa, H. P. and Rousseeuw, P. J. (1991), “Breakdown points of affine equivariant estimators of multivariate location and covariance matrices,” *The Annals of Statistics*, 229–248.
- Markopoulos, P. P., Karystinos, G. N., and Pados, D. A. (2014), “Optimal algorithms for L_1 -subspace signal processing,” *IEEE Transactions on Signal Processing*, 62, 5046–5058.
- Prothero, J., Jiang, M., Hannig, J., Tran-Dinh, Q., Ackerman, A., and Marron, J. (2022), “Data integration via analysis of subspaces (DIVAS),” *arXiv preprint arXiv:2212.00703*.
- Rahmani, M. and Atia, G. K. (2017), “Coherence pursuit: Fast, simple, and robust principal component analysis,” *IEEE Transactions on Signal Processing*, 65, 6260–6275.
- Rousseeuw, P. J. (1984), “Least median of squares regression,” *Journal of the American Statistical Association*, 79, 871–880.
- (1985), “Multivariate estimation with high breakdown point,” *Mathematical statistics and applications*, 8, 37.
- Shamir, O. (2016), “Fast stochastic algorithms for SVD and PCA: Convergence properties and convexity,” in *International Conference on Machine Learning*, PMLR, pp. 248–256.
- She, Y., Li, S., and Wu, D. (2016), “Robust orthogonal complement principal component analysis,” *Journal of the American Statistical Association*, 111, 763–771.
- Sykulski, M. (2015), *rpca: RobustPCA: Decompose a Matrix into Low-Rank and Sparse Components*, r package version 0.2.3.
- Tang, P. and Phillips, J. M. (2016), “The robustness of estimator composition,” *Advances in Neural Information Processing Systems*, 29.
- Tyler, D. E., Yi, M., and Nordhausen, K. (2023), “Robust and Resistant Regularized Covariance Matrices,” *arXiv preprint arXiv:2307.15774*.
- Vershynin, R. (2018), *High-dimensional probability: An introduction with applications in data science*, vol. 47, Cambridge university press.
- Wright, J., Ganesh, A., Min, K., and Ma, Y. (2013), “Compressive principal component pursuit,” *Information and Inference: A Journal of the IMA*, 2, 32–68.
- Xu, C., Xu, W., and Jing, K. (2023), “Fast algorithms for singular value decomposition and inverse of nearly low-rank matrices,” *National Science Review*, nwad083.
- Xu, H., Caramanis, C., and Mannor, S. (2012), “Outlier-robust PCA: The high-dimensional case,” *IEEE transactions on information theory*, 59, 546–572.
- Xu, H., Caramanis, C., and Sanghavi, S. (2010), “Robust PCA via outlier pursuit,” *Advances in Neural Information Processing Systems*, 23.
- Yi, X., Park, D., Chen, Y., and Caramanis, C. (2016), “Fast algorithms for robust PCA via gradient descent,” *Advances in Neural Information Processing Systems*, 29.
- Zhang, L., Shen, H., and Huang, J. Z. (2013), “Robust regularized singular value decomposition with application to mortality data,” *The Annals of Applied Statistics*, 1540–1561.
- Zhou, Z., Li, X., Wright, J., Candes, E., and Ma, Y. (2010), “Stable principal component pursuit,” in *2010 IEEE International Symposium on Information Theory*, IEEE, pp. 1518–1522.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]

3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplementary material or as a URL). [No]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [No]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Yes]
 - (c) New assets either in the supplementary material or as a URL, if applicable. [No]
 - (d) Information about consent from data providers/curators. [Yes]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

7 Supplementary material

7.1 TECHNICAL DETAILS

7.1.1 Statistical Accuracy Theorem

A statistical accuracy of \mathbf{v}_r^{SP} can be measured asymptotically. For this purpose, assume that each p -vector \mathbf{x}_i is independently sampled from a mean-zero elliptical distribution (Cambanis et al., 1981), \mathcal{F}_Σ with covariance matrix $\Sigma = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T = \sum_{i=1}^p \lambda_i \mathbf{v}_i \mathbf{v}_i^T$ (satisfying $\lambda_i > \lambda_{i+1}$). Translating a result of Graciela and Fraiman (1999) on the robust PCA estimation of Locantore et al. (1999), we observe the following:

Theorem (Graciela and Fraiman (1999)). *Let $\mathbf{X}_n = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T$ for increasing n , where \mathbf{x}_i 's are independently sampled from \mathcal{F}_Σ . If $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_p)$ consists of distinct non-negative diagonal elements, then for $r = 1, \dots, p$, $\mathbf{v}_r^{\text{SP}}(\mathbf{X}_n)$ is a consistent estimator of the r th PC direction \mathbf{v}_r in the sense that $\|\mathbf{v}_r - \mathbf{v}_r^{\text{SP}}(\mathbf{X}_n)\|_2 \rightarrow 0$ almost surely as $n \rightarrow \infty$.*

Note that the conclusion of above theorem holds when $\mathbf{v}_r^{\text{SP}}(\mathbf{X}_n)$ is replaced by the right singular vector $\mathbf{v}_r(\mathbf{X}_n)$, which may be used to justify the use of SVD in the estimation of PC directions. The theorem implies that $\mathbf{v}_r^{\text{SP}}(\mathbf{X})$ is a statistically accurate approximation of the population PC direction.

Proof. It is known that, for independent and identically distributed (i.i.d.) random variables with a finite covariance matrix, the sample covariance matrix almost surely converges to the population covariance matrix. Furthermore, when the eigenvalues of the population covariance matrix are distinct, the eigenvector corresponding to the j th largest eigenvalue of the sample covariance matrix converges to that of the population covariance matrix, almost surely.

It is enough to show that $\text{Var}(\frac{\mathbf{x}_1}{\|\mathbf{x}_1\|_2}) = \mathbf{V}\tilde{\mathbf{\Lambda}}\mathbf{V}^T$ where $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_p)$ is a matrix consisting of PC directions, and $\tilde{\mathbf{\Lambda}} = \text{Diag}(\tilde{\lambda}_1, \dots, \tilde{\lambda}_p)$ with $\tilde{\lambda}_1 > \dots > \tilde{\lambda}_p$. Denote $\mathbf{Z} = \mathbf{V}^T \mathbf{x}_1$. Note that the \mathbf{x}_1 is from an elliptical distribution, so that characteristic function of \mathbf{x}_1 is represented by $\phi_{\mathbf{x}_1}(\mathbf{t}) = \psi(\mathbf{t}^T \mathbf{V}\mathbf{\Lambda}\mathbf{V}\mathbf{t})$ for some scalar function $\psi: \mathbb{R} \rightarrow \mathbb{R}$, and the characteristic function of \mathbf{Z} is represented by $\phi_{\mathbf{Z}}(\mathbf{t}) = \psi(\mathbf{t}^T \mathbf{\Lambda}\mathbf{t})$. It implies that $\text{pdf}_{\mathbf{Z}}(z_1, \dots, z_j, \dots, z_p) = \text{pdf}_{\mathbf{Z}}(z_1, \dots, -z_j, \dots, z_p)$ for any $j = 1, \dots, p$. Thus, we have

$$\mathbb{E}\left(\frac{Z_i}{\|\mathbf{Z}\|_2}\right) = 0, \quad \mathbb{E}\left(\frac{Z_i Z_j}{\|\mathbf{Z}\|_2^2}\right) = 0.$$

It implies that the off diagonal elements of $\text{Var}(\frac{\mathbf{Z}}{\|\mathbf{Z}\|_2})$ are zero. On the other, for $j = 1, \dots, p-1$,

$$\begin{aligned} \mathbb{E}\left(\frac{Z_j^2}{\|\mathbf{Z}\|_2^2}\right) &= \mathbb{E}\left(\frac{Z_j^2}{Z_j^2 + Z_{j+1}^2 + \sum_{i \neq j, j+1} Z_i^2}\right) \\ &= \mathbb{E}\left(\frac{\lambda_j Y_j^2}{\lambda_j Y_j^2 + \lambda_{j+1} Y_{j+1}^2 + \sum_{i \neq j, j+1} \lambda_i Y_i^2}\right) \\ &> \mathbb{E}\left(\frac{\lambda_{j+1} Y_{j+1}^2}{\lambda_{j+1} Y_{j+1}^2 + \lambda_j Y_j^2 + \sum_{i \neq j, j+1} \lambda_i Y_i^2}\right) = \mathbb{E}\left(\frac{Z_{j+1}^2}{\|\mathbf{Z}\|_2^2}\right) \end{aligned}$$

where $\mathbf{Y} = (Y_1, \dots, Y_p)^T$ is a random vector which has a characteristic function $\phi_{\mathbf{Y}}(\mathbf{t}) = \psi(\mathbf{t}^T \mathbf{t})$. Thus, the diagonal elements of $\text{Var}(\frac{\mathbf{Z}}{\|\mathbf{Z}\|_2})$ are decreasing. Hence,

$$\text{Var}\left(\frac{\mathbf{x}_1}{\|\mathbf{x}_1\|_2}\right) = \mathbf{V}\text{Var}\left(\frac{\mathbf{Z}}{\|\mathbf{Z}\|_2}\right)\mathbf{V}^T = \mathbf{V}\tilde{\mathbf{\Lambda}}\mathbf{V}^T$$

where $\tilde{\mathbf{\Lambda}} = \text{Diag}(\tilde{\lambda}_1, \dots, \tilde{\lambda}_p)$ with $\tilde{\lambda}_j = \mathbb{E}\left(\frac{Z_j^2}{\|\mathbf{Z}\|_2^2}\right)$. □

7.1.2 Proof of Theorem 1

Proof. Let $\Sigma_0 = \sum_{j=1}^p d_j \mathbf{v}_j \mathbf{v}_j^T$ be the covariance matrix of $\frac{\mathbf{x}_1}{\|\mathbf{x}_1\|_2}$ by Statistical Accuracy Theorem. Let $\hat{\Sigma}_\epsilon$ be the sample covariance made with n normalized samples $[\frac{\mathbf{x}_1}{\|\mathbf{x}_1\|}, \dots, \frac{\mathbf{x}_{(1-\epsilon)n}}{\|\mathbf{x}_{(1-\epsilon)n}\|}, \frac{\mathbf{y}_1}{\|\mathbf{y}_1\|}, \dots, \frac{\mathbf{y}_{\epsilon n}}{\|\mathbf{y}_{\epsilon n}\|}]$. By Davis-Kahan

theorem and taking expectation the both sides, we have $\mathbb{E} \left[\min\{\|\hat{\mathbf{v}}_j^{\text{Sp}} - \mathbf{v}_j\|_2, \|\hat{\mathbf{v}}_j^{\text{Sp}} + \mathbf{v}_j\|_2\} \right] \leq \frac{2^{2/3}}{\delta_j} \mathbb{E} \left[\|\hat{\boldsymbol{\Sigma}}_\epsilon - \boldsymbol{\Sigma}_0\| \right]$ with spectral norm $\|\cdot\|$. The right-hand side can be bounded as

$$\mathbb{E} \left[\|\hat{\boldsymbol{\Sigma}}_\epsilon - \boldsymbol{\Sigma}_0\| \right] \leq \epsilon \mathbb{E} \left[\left\| \frac{1}{\epsilon n} \sum_{i=1}^{\epsilon n} \frac{\mathbf{y}_i \mathbf{y}_i^T}{\|\mathbf{y}_i\|_2^2} \right\| \right] + (1 - \epsilon) \mathbb{E} \left[\left\| \frac{1}{(1 - \epsilon)n} \sum_{i=1}^{(1 - \epsilon)n} \frac{\mathbf{x}_i \mathbf{x}_i^T}{\|\mathbf{x}_i\|_2^2} - \boldsymbol{\Sigma}_0 \right\| \right].$$

Using Theorem 4.7.1 (Covariance estimation in Vershynin (2018)), We have

$$\mathbb{E} \left[\|\hat{\boldsymbol{\Sigma}}_\epsilon - \boldsymbol{\Sigma}_0\| \right] \leq \epsilon + (1 - \epsilon) c \left(\sqrt{\frac{p}{(1 - \epsilon)n}} + \frac{p}{(1 - \epsilon)n} \right),$$

where c is an absolute constant. Thus,

$$\begin{aligned} \mathbb{E} \left[\|\hat{\mathbf{v}}_j^{\text{Sp}} - \mathbf{v}_j\|_2 \right] &\leq \frac{1}{\delta_j} \left\{ 2^{\frac{3}{2}} \epsilon + (1 - \epsilon) C \left(\sqrt{\frac{p}{(1 - \epsilon)n}} + \frac{p}{(1 - \epsilon)n} \right) \right\} \\ &\leq \frac{1}{\delta_j} (C' \epsilon + C'' \sqrt{(1 - \epsilon)\epsilon}) \end{aligned}$$

for some absolute constants $C', C'' > 0$. □

7.1.3 Proof of Lemma 2

We provide a proof for the lemma, where Lemma 2 is a direct consequence of the following lemma.

Lemma 6.

- (i) $\text{bp}_{\text{row}}(\mathcal{V}; \mathbf{X}) = k$ if and only if $(k, l) \in \text{BP}(\mathcal{V}; \mathbf{X})$ for some $1 \leq l \leq p$ and $(i, j) \notin \text{BP}(\mathcal{V}; \mathbf{X})$ for any $(i, j) \preceq (k - 1, p)$.
- (ii) $\text{bp}_{\text{col}}(\mathcal{V}; \mathbf{X}) = l$ if and only if $(k, l) \in \text{BP}(\mathcal{V}; \mathbf{X})$ for some $1 \leq k \leq n$ and $(i, j) \notin \text{BP}(\mathcal{V}; \mathbf{X})$ for any $(i, j) \preceq (n, l - 1)$.

Proof. Assume that $\text{bp}_{\text{row}}(\mathcal{V}; \mathbf{X}) = k$. It implies that \mathcal{V} breaks down at block-size (k, p) at \mathbf{X} , thus there exists $(i, j) \in \text{BP}(\mathcal{V}; \mathbf{X})$ such that $(i, j) \preceq (k, p)$. Since $\text{bp}_{\text{row}}(\mathcal{V}; \mathbf{X}) > k - 1$, \mathcal{V} does not break down at block-size (i, j) at \mathbf{X} for any $(i, j) \preceq (k - 1, p)$. Thus, $(k, l) \in \text{BP}(\mathcal{V}; \mathbf{X})$ for some $1 \leq l \leq p$ and $(i, j) \notin \text{BP}(\mathcal{V}; \mathbf{X})$ for any $(i, j) \preceq (k - 1, p)$.

Assume that $(k, l) \in \text{BP}(\mathcal{V}; \mathbf{X})$ for some $1 \leq l \leq p$ and $(i, j) \notin \text{BP}(\mathcal{V}; \mathbf{X})$ for any $(i, j) \preceq (k - 1, p)$. Then, we have $\text{bp}_{\text{row}}(\mathcal{V}; \mathbf{X}) \leq k$ since $(k, l) \in \text{BP}(\mathcal{V}; \mathbf{X})$ for some $1 \leq l \leq p$, and $\text{bp}_{\text{row}}(\mathcal{V}; \mathbf{X}) > k - 1$ since $(i, j) \notin \text{BP}(\mathcal{V}; \mathbf{X})$ for any $(i, j) \preceq (k - 1, p)$.

The proof of the second part of this lemma about the column-wise breakdown points, can be given by following the lines of the above, with \mathbf{X} replaced by \mathbf{X}^T and \mathcal{V} replaced by $\mathcal{V}^T : \mathbb{R}^{p \times n} \rightarrow \mathbb{R}$ given by $\mathcal{V}^T(\mathbf{X}^T) = \mathcal{V}(\mathbf{X})$. Lemma 2 is proved by setting $k = 1$ or $l = 1$ in Lemma 6. □

7.1.4 Proof of Proposition 3

Proof. The proof of Proposition 3 can be given by following the lines of the proof for Theorem 4, with $\rho(\mathbf{Z}) = \sum_{i,j} \rho_{ij}(z_{ij}) = \sum_{i,j} z_{ij}^2$ without the regularizations. □

7.1.5 Proof of Theorem 4

Proof. Let $(\sigma_r^\rho, \mathbf{u}_r^\rho, \mathbf{v}_r^\rho)$ be the r th solution to (7) for $\mathbf{X} = (x_{ij}) \in \mathbb{R}^{n \times p}$. Let $\mathcal{U} = \mathcal{U}_R^\rho(\mathbf{X})$ and let \mathcal{U}^\perp be the orthogonal complement of $\mathcal{U}_R^\rho(\mathbf{X})$. Denote a unit vector in \mathbb{R}^n with the i th element 1 by \mathbf{e}_i . Since the dimension of \mathcal{U} is $R < n$, there exist $\mathbf{w} := d_1 \mathbf{e}_{l_1} + \dots + d_{R+1} \mathbf{e}_{l_{R+1}}$ with $d_1^2 + \dots + d_{R+1}^2 = 1$ such that $\mathbf{w} \in \mathcal{U}^\perp$. Assume that $l_i = i$, for notational simplicity and denote $\mathbf{w} = (w_1, \dots, w_n)^T = (d_1, \dots, d_{R+1}, 0, \dots, 0)^T$.

Let M be a constant satisfying $\sum_{(i,j) \in I} \rho_{ij}(x_{ij}) + b < M$ where $I = \{(i,j) : i > R+1 \text{ or } j > 1\}$ and $b = \sup_{\mathbf{u} \in \mathcal{S}^{n-1}} \mathcal{P}_1(\mathbf{u}) + \sup_{\mathbf{v} \in \mathcal{S}^{p-1}} \mathcal{P}_2(\mathbf{v})$. For arbitrary $0 < \varepsilon < 1$, define $\mathcal{S}_\varepsilon = \{\mathbf{u} \in \mathbb{R}^n : |\mathbf{u}^T \mathbf{w}| > \|\mathbf{u}\|_2(1 - \varepsilon^2/2)\}$ and $\mathcal{B}(c) = \{\mathbf{u} = (u_1, \dots, u_n) \in \mathbb{R}^n : \sum_{i=1}^{R+1} \rho_{i1}(cd_i - u_i) + \sum_{i=R+2}^n \rho_{i1}(x_{i1} - u_i) < M\}$. For some $M' > 0$, we can find a neighborhood of $c\mathbf{w}$ with radius M' , given by $B(c\mathbf{w}, M') := \{c\mathbf{w} + \mathbf{u} \in \mathbb{R}^n : \|\mathbf{u}\|_2 \leq M'\}$ satisfying $\mathcal{B}(c) \subset B(c\mathbf{w}, M')$ since ρ_{ij} diverges at the boundaries of \mathbb{R} . Here $M' > 0$ is independent of $c > 0$. By growing c , we can choose c_0 such that $\mathcal{B}(c_0) \subset B(c_0\mathbf{w}, M') \subset \mathcal{S}_\varepsilon$. Let $\mathbf{Z} = (z_{ij}) = (\mathbf{z}^1, \dots, \mathbf{z}^p) \in \mathbb{R}^{n \times p}$ be denoted by

$$\begin{cases} z_{ij} = x_{ij} & \text{if } (i,j) \in I, \\ z_{i1} = c_0 d_i & \text{if } i = 1, \dots, R+1. \end{cases}$$

Then,

$$\begin{aligned} & \min_{\mathbf{a} \in \mathbb{R}^p} \{\rho(\mathbf{Z} - \mathbf{w}\mathbf{a}^T) + \mathcal{P}_1(\mathbf{w}) + \mathcal{P}_2(\mathbf{a}/\|\mathbf{a}\|)\} \\ & \leq \min_{a_1 \in \mathbb{R}} \sum_{i=1}^{R+1} \rho_{i1}(z_{i1} - d_i a_1) + \sum_{j=2}^p \min_{a_j \in \mathbb{R}} \sum_{i=1}^{R+1} \rho_{ij}(z_{ij} - d_i a_j) + \sum_{i=R+2}^n \sum_{j=1}^p \rho_{ij}(z_{ij}) + b \\ & = \min_{a_1 \in \mathbb{R}} \sum_{i=1}^{R+1} \rho_{i1}(c_0 d_i - d_i a_1) + \sum_{j=2}^p \min_{a_j \in \mathbb{R}} \sum_{i=1}^{R+1} \rho_{ij}(x_{ij} - d_i a_j) + \sum_{i=R+2}^n \sum_{j=1}^p \rho_{ij}(x_{ij}) + b \\ & \leq 0 + \sum_{(i,j) \in I} \rho_{ij}(x_{ij}) + b < M. \end{aligned}$$

For any $\mathbf{u} \notin \mathcal{S}_\varepsilon$ with $\|\mathbf{u}\|_2 = 1$, $\mathbf{a}\mathbf{u} \notin \mathcal{B}(c_0)$ for any $a \in \mathbb{R}$. Hence,

$$\begin{aligned} & \min_{\mathbf{a} \in \mathbb{R}^p} \{\rho(\mathbf{Z} - \mathbf{u}\mathbf{a}^T) + \mathcal{P}_1(\mathbf{w}) + \mathcal{P}_2(\mathbf{a}/\|\mathbf{a}\|)\} \\ & \geq \min_{a_1 \in \mathbb{R}} \sum_{i=1}^n \rho_{i1}(z_{i1} - v_i a_1) \geq M. \end{aligned}$$

Thus, $\mathbf{w}_1^\rho \in \mathcal{S}_\varepsilon$ where $(\boldsymbol{\eta}_r^\rho, \mathbf{w}_r^\rho, \mathbf{s}_r^\rho)$ is the r th solution to (7) for \mathbf{Z} . Hence we have $|(\mathbf{w}_1^\rho)^T \mathbf{u}_r^\rho| = |(\mathbf{w}_1^\rho - \mathbf{w} + \mathbf{w})^T \mathbf{u}_r^\rho| \leq \sqrt{2 - 2(1 - \varepsilon^2/2)} = \varepsilon$ for all r . Let $\mathbf{U}^\rho = (\mathbf{u}_1^\rho, \dots, \mathbf{u}_R^\rho) \in \mathbb{R}^{n \times R}$ be a matrix whose columns are basis vectors of \mathcal{U} , and $\mathbf{W}^\rho = (\mathbf{w}_1^\rho, \dots, \mathbf{w}_R^\rho) \in \mathbb{R}^{n \times R}$ be a matrix whose columns are basis vectors of $\mathcal{W} := \mathcal{U}_R^\rho(\mathbf{Z})$. Let $\mathbf{A} = \mathbf{W}^T \mathbf{V} \mathbf{V}^T \mathbf{W}$ and $\tilde{\mathbf{A}} = (\tilde{a}_{ij})$ with the (i,j) th element $\tilde{a}_{ij} = 0$ if $i = 1$ or $j = 1$, and $\tilde{a}_{ij} = a_{ij}$ otherwise. By Weyl's Theorem,

$$\cos^2(\theta(\mathcal{U}, \mathcal{W})) \leq \|\mathbf{A} - \tilde{\mathbf{A}}\|_F \leq 2R^2 \varepsilon.$$

Since ε is arbitrary, $\sup_{\mathbf{Z}_{R+1}^1} \theta(\mathcal{U}_R(\mathbf{X}), \mathcal{U}_R(\mathbf{Z}_{R+1}^1)) = \frac{\pi}{2}$. Thus, $(k, 1) \in \text{BP}(\mathcal{U}_R^\rho; \mathbf{X})$ for some $k \leq R+1$. Moreover, it implies that, $\text{bp}_{\text{col}}(\mathcal{U}_R^\rho; \mathbf{X}) = 1$ and $\text{bp}_{\text{row}}(\mathcal{U}_R^\rho; \mathbf{X}) \leq R+1$.

The proof of the second part of this theorem can be given by following the lines of the above, with \mathbf{X} replaced by \mathbf{X}^T and \mathcal{U}_R^ρ replaced by $\mathcal{U}_R^{\rho T} : \mathbb{R}^{p \times n} \rightarrow \mathbb{R}$ given by $\mathcal{U}_R^{\rho T}(\mathbf{X}^T) = \mathcal{V}_R^\rho(\mathbf{X})$. \square

7.1.6 Proof of Theorem 5

Proof. We prove the second part of this theorem about the R th right singular space, $\mathcal{V}_R^{\text{Sp}}$. Assume that $\text{bp}_{\text{row}}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X}) = m < n_R$. For an arbitrary small ε , take a corrupted data $\mathbf{Z}_m = (\mathbf{z}_1, \dots, \mathbf{z}_n)$ such that $|I_0| = |\{i : \mathbf{z}_i = \mathbf{x}_i\}| = n - m$, and choose $\mathbf{v}_\perp \in \mathcal{V}_R^{\text{Sp}}(\mathbf{Z}_m)$ such that $\|\mathbf{\Pi}_R^\perp \mathbf{v}_\perp\|_2 < \varepsilon$ where $\mathbf{\Pi}_R^\perp$ is the projection

matrix of $\mathcal{V}_R^{\text{Sp}}(\mathbf{X})$. Let $\mathbf{\Pi}_R = (\mathbf{I} - \mathbf{\Pi}_R^\perp)$ be the projection matrix of the orthogonal complement of $\mathcal{V}_R^{\text{Sp}}(\mathbf{X})$. Then,

$$\begin{aligned} \mathbf{v}_\perp^T \left(\sum_{i=1}^n \frac{\mathbf{z}_i \mathbf{z}_i^T}{\mathbf{z}_i^T \mathbf{z}_i} \right) \mathbf{v}_\perp &= \mathbf{v}_\perp^T \left(\sum_{i \notin I_0} \frac{\mathbf{z}_i \mathbf{z}_i^T}{\mathbf{z}_i^T \mathbf{z}_i} \right) \mathbf{v}_\perp + \mathbf{v}_\perp^T \left(\sum_{i \in I_0} \frac{\mathbf{x}_i \mathbf{x}_i^T}{\mathbf{x}_i^T \mathbf{x}_i} \right) \mathbf{v}_\perp \\ &\leq m + \mathbf{v}_\perp^T \left(\sum_{i \in I_0} (\mathbf{\Pi}_R + \mathbf{\Pi}_R^\perp) \frac{\mathbf{x}_i \mathbf{x}_i^T}{\mathbf{x}_i^T \mathbf{x}_i} (\mathbf{\Pi}_R + \mathbf{\Pi}_R^\perp) \right) \mathbf{v}_\perp \\ &\leq m + \mathbf{v}_\perp^T \left(\sum_{i \in I_0} \mathbf{\Pi}_R \frac{\mathbf{x}_i \mathbf{x}_i^T}{\mathbf{x}_i^T \mathbf{x}_i} \mathbf{\Pi}_R \right) \mathbf{v}_\perp + 3(n-m)\varepsilon \\ &\leq m + (\lambda_1(\tilde{\mathbf{X}}_m \mathbf{\Pi}_R))^2 + 3(n-m)\varepsilon \end{aligned}$$

where $\tilde{\mathbf{X}}_m \in \mathbb{R}^{(n-m) \times p}$ is the submatrix of $\tilde{\mathbf{X}}_{\text{row}}$ obtained by choosing $(n-m)$ rows indexed by I_0 . On the other hand,

$$\begin{aligned} \mathbf{v}_\perp^T \left(\sum_{i=1}^n \frac{\mathbf{z}_i \mathbf{z}_i^T}{\mathbf{z}_i^T \mathbf{z}_i} \right) \mathbf{v}_\perp &\geq \sup_{\mathbf{v} \in \text{Gr}(p, R)} \min_{\mathbf{v} \in \mathcal{V}} \left\{ \mathbf{v}^T \left(\sum_{i=1}^n \frac{\mathbf{z}_i \mathbf{z}_i^T}{\mathbf{z}_i^T \mathbf{z}_i} \right) \mathbf{v} \right\} \\ &= \sup_{\mathbf{v} \in \text{Gr}(p, R)} \min_{\mathbf{v} \in \mathcal{V}} \left\{ \mathbf{v}^T \left(\sum_{i \notin I_0} \frac{\mathbf{z}_i \mathbf{z}_i^T}{\mathbf{z}_i^T \mathbf{z}_i} \right) \mathbf{v} + \mathbf{v}^T \left(\sum_{i \in I_0} \frac{\mathbf{x}_i \mathbf{x}_i^T}{\mathbf{x}_i^T \mathbf{x}_i} \right) \mathbf{v} \right\} \\ &\geq (\lambda_R(\tilde{\mathbf{X}}_m))^2. \end{aligned}$$

It implies that

$$(1 - 3\varepsilon)m \geq (\lambda_R(\tilde{\mathbf{X}}_m))^2 - (\lambda_1(\tilde{\mathbf{X}}_m \mathbf{\Pi}_R))^2 - 3n\varepsilon,$$

thus,

$$(1 - 3\varepsilon)m \geq \inf_{\tilde{\mathbf{X}}_m} \{ (\lambda_R(\tilde{\mathbf{X}}_m))^2 - (\lambda_1(\tilde{\mathbf{X}}_m \mathbf{\Pi}_R))^2 \} - 3n\varepsilon.$$

Since ε is arbitrary,

$$m \geq \inf_{\tilde{\mathbf{X}}_m} \{ (\lambda_R(\tilde{\mathbf{X}}_m))^2 - (\lambda_1(\tilde{\mathbf{X}}_m \mathbf{\Pi}_R))^2 \}.$$

It is a contradiction since $m < n_R$. Thus, $\text{bp}_{\text{row}}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X}) \geq n_R$. It implies that $(n_R, 1) \preceq (i, j)$ for any $(i, j) \in \text{BP}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X})$.

To show that $\text{bp}_{\text{col}}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X}) \leq R + 1$, let $\mathcal{V} = \mathcal{V}_R^{\text{Sp}}(\mathbf{X})$ and let \mathcal{V}^\perp be the orthogonal complement of $\mathcal{V}_R^{\text{Sp}}(\mathbf{X})$. Denote the unit vector in \mathbb{R}^p with the j th element 1 by \mathbf{e}^j . Since the dimension of \mathcal{V} is $R < p$, there exist $\mathbf{s} := d_1 \mathbf{e}^{l_1} + \dots + d_{R+1} \mathbf{e}^{l_{R+1}}$ with $d_1^2 + \dots + d_{R+1}^2 = 1$ such that $\mathbf{s} \in \mathcal{V}^\perp$. Assume that $l_j = j$, for notational simplicity and denote $\mathbf{s} = (w_1, \dots, w_p)^T = (d_1, \dots, d_{R+1}, 0, \dots, 0)^T$. For $\mathbf{X} = (\mathbf{x}^1, \dots, \mathbf{x}^p) \in \mathbb{R}^{n \times p}$ and $c > 0$, define a corrupted matrix constructed by replacing $R + 1$ columns of \mathbf{X} , by $\mathbf{Z}^{R+1}(c) = (\mathbf{z}_1(c), \dots, \mathbf{z}_n(c))^T = (\mathbf{z}^1(c), \dots, \mathbf{z}^p(c))$ with $\mathbf{z}^j(c) = cd_j \mathbf{1}_n$ for $j = 1, \dots, R + 1$, and $\mathbf{z}^j(c) = \mathbf{x}^j$ for $j = R + 2, \dots, p$. Let $\mathbf{Y}(c) = (\mathbf{y}_1(c), \dots, \mathbf{y}_n(c))^T \in \mathbb{R}^{n \times p}$ be a row-normalized matrix of $\mathbf{Z}^{R+1}(c)$ whose i th row is $\mathbf{y}_i(c) = \frac{\mathbf{z}_i(c)}{\|\mathbf{z}_i(c)\|_2}$. Denote $\mathbf{Y} := \lim_{c \rightarrow \infty} \mathbf{Y}(c) = (\mathbf{s}, \dots, \mathbf{s})^T \in \mathbb{R}^{n \times p}$. By Davis-Khan Theorem,

$$\sin \left(\theta(\mathbf{s}, \mathbf{v}_1(\mathbf{Y}(c))) \right) \leq \frac{2}{n} \|\mathbf{Y}^T \mathbf{Y} - \mathbf{Y}(c)^T \mathbf{Y}(c)\|_{\text{F}} \rightarrow 0$$

as $c \rightarrow \infty$ where $\mathbf{v}_1(\mathbf{Y}(c))$ is the right singular vector corresponding to the largest singular value of $\mathbf{Y}(c)$. Note that the right singular space obtained from SpSVD, $\mathcal{V}_R^{\text{Sp}}(\mathbf{Z}(c))$, equals the right singular space, $\mathcal{V}_R(\mathbf{Y}(c))$. Thus

$$\begin{aligned} \sup_{\mathbf{Z}^{R+1}} \theta(\mathcal{V}_R^{\text{Sp}}(\mathbf{X}), \mathcal{V}_R^{\text{Sp}}(\mathbf{Z}^{R+1})) &\geq \lim_{c \rightarrow \infty} \theta(\mathcal{V}_R^{\text{Sp}}(\mathbf{X}), \mathcal{V}_R^{\text{Sp}}(\mathbf{Z}^{R+1}(c))) \\ &= \lim_{c \rightarrow \infty} \theta(\mathcal{V}_R^{\text{Sp}}(\mathbf{X}), \mathcal{V}_R(\mathbf{Y}(c))) = \frac{\pi}{2}. \end{aligned}$$

It implies that $\text{bp}_{\text{col}}(\mathcal{V}_R^{\text{Sp}}; \mathbf{X}) \leq R + 1$.

The proof of the first part of this theorem can be given by following the lines of the above, with \mathbf{X} replaced by \mathbf{X}^T and $\mathcal{V}_R^{\text{Sp}}$ replaced by $(\mathcal{V}_R^{\text{Sp}})^T : \mathbb{R}^{p \times n} \rightarrow \mathbb{R}$ given by $(\mathcal{V}_R^{\text{Sp}})^T(\mathbf{X}^T) = \mathcal{U}_R^{\text{Sp}}(\mathbf{X})$. \square

7.2 COMPUTATION COMPLEXITY

7.2.1 Computation Complexity for Algorithm 1

For inputs, small rank R and a data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, Algorithm 1 consists of the following steps: normalization, rank- R SVD, and finding the solution to (3). Each row of \mathbf{X} , which is a p -vector, is normalized for $i = 1, \dots, n$. This step requires computation time of $O(np)$. Similarly, column normalization is performed with the same complexity. We use a partial SVD algorithm to find the rank- R approximation of the normalized data, which requires $O(npR)$ complexity. Finally, the solution to (3) is computed using a selection algorithm that finds the smallest value among the npR^2 elements. Each step for $r = 1, \dots, R$ takes npR^2 computations, resulting in a total computation burden of $O(npR^3)$ for $k = 1, \dots, R$.

7.2.2 Computation Complexity of Other Methods

In Section 5, we performed comparison of SpSVD with other methods. Brahma et al. (2017) presents computational complexity of part of the R2PCP algorithm in their paper. Although they propose new batch version of algorithm to get reduced computational complexity, they did not reveal the total computational complexity of the entire process. On the other hand, Candès et al. (2011) claims that the dominant cost in RPCA algorithm comes from computing one partial SVD per iteration. Although the number of iterations in their algorithm appears to remain nearly constant regardless of dimension, the algorithm in the R package `rpca` (Sykulski, 2015) requires significantly more time to run compared to other methods.

7.3 ADDITIONAL DETAILS IN NUMERICAL STUDIES

In this numerical studies, we generate the data, and implement standard SVD and SpSVD using R. We used the R package `rpca` (Sykulski, 2015) for RPCA, `Matlab` implementation for COP, and R2PCP (Brahma et al., 2017) was implemented using their source code in `Matlab`.

7.3.1 Data Generation and Adding Contamination in the Simulation Experiment

To generate randomly sampled matrices \mathbf{U} and \mathbf{V} from the uniform distribution on the set of orthogonal matrices, respectively, we create a matrix consisting of standard normal random noise and apply QR decomposition. It is well known that this orthogonal matrix obtained from QR decomposition follows the uniform distribution on the set of orthogonal matrices of the same size.

To generate the sparse outlier matrix \mathbf{S} , we arbitrary choose 0.05% of the indices as $I = \{i_1, \dots, i_{10}\} \subset \{1, \dots, n\}$ and $J = \{j_1, \dots, j_5\} \subset \{1, \dots, p\}$, respectively. Note that $|I| \geq 4$ and $|J| \geq 4$. Let $\mathbf{U}_I \in \mathbb{R}^{|I| \times 3}$ and $\mathbf{V}_J \in \mathbb{R}^{|J| \times 3}$ be the submatrices of $\mathbf{U} \in \mathbb{R}^{n \times 3}$ and $\mathbf{V} \in \mathbb{R}^{p \times 3}$ consisting of rows indexed by I and J , respectively. We can find solutions $\tilde{\mathbf{a}} = (\tilde{a}_1, \dots, \tilde{a}_{10})' \in \mathbb{R}^{|I|}$ and $\tilde{\mathbf{b}} = (\tilde{b}_1, \dots, \tilde{b}_5)' \in \mathbb{R}^{|J|}$ satisfying $\mathbf{U}_I^T \tilde{\mathbf{a}} = \mathbf{0}_3$ and $\mathbf{V}_J^T \tilde{\mathbf{b}} = \mathbf{0}_3$. We define the sparse outlier matrix by $\mathbf{S} = \mathbf{a}\mathbf{b}^T / \|\mathbf{a}\|_2 \|\mathbf{b}\|_2$, where $\mathbf{a} = (a_1, \dots, a_n)' \in \mathbb{R}^n$ has elements zero excepts for elements indexed by I as $a_{i_l} = \tilde{a}_l$ for $l = 1, \dots, 10$, and $\mathbf{b} = (b_1, \dots, b_p)' \in \mathbb{R}^p$ has elements zero excepts for elements indexed by J as $b_{j_k} = \tilde{b}_k$ for $k = 1, \dots, 5$. Then, the row space and column space of \mathbf{S} are orthogonal to those spaces of \mathbf{L} , respectively.

7.3.2 Extra Result for the Simulation Experiment

Due to computational limitations, we compared SVD, SpSVD, and RPCA without iterations for increasing matrix sizes in Section 5. In addition to those comparisons, we also provide experiments with 100 iterations while increasing the matrix size from 200×100 to 600×300 . For $C = 1, 1.2, \dots, 2.8, 3$, with $n = 200$ and $p = 100$, we generate a data matrix $\mathbf{X} \in \mathbb{R}^{C \cdot n \times C \cdot p}$. To maintain the magnitude of singular values and outliers, we set η as $500C$ and the singular values (d_1, d_2, d_3) as $C \cdot (80, 70, 60)$. The block of \mathbf{S} containing the outliers has sizes of $(0.05nC, 0.05pC)$. The results are presented in Figure 2.

As C increases, RPCA and SpSVD exhibit better accuracy for the singular spaces and a singular value in Figure 2(a), 2(b), and 2(c). RPCA takes longer time as C increases compared to SVD, and other methods including ELSVD and R2PCP also require more time as C increases compared to SVD as described in Figure 2(d).

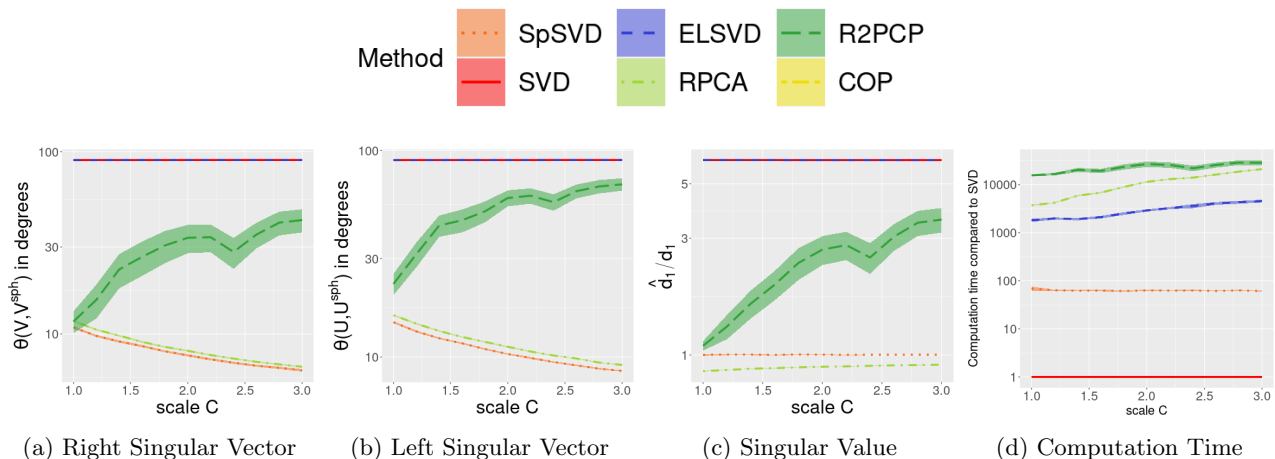


Figure 2: The Approximation Accuracy, Robustness, and Computation Times against Increasing Scale C in (a), (b), (c) and (d).

7.3.3 Higher Rank Case

The data matrix without outliers $\mathbf{X} \in \mathbb{R}^{n \times p}$ is set to be the sum of rank-9 $\mathbf{L} = (l_{ij})$ and \mathbf{E} consisting of standard normal random noises with $n = 1000$ and $p = 500$. The rank-9 \mathbf{L} is $\mathbf{L} = \sum_{r=1}^9 d_r \mathbf{u}_r \mathbf{v}_r^T$, $(d_1, d_2, \dots, d_8, d_9) = (750, 700, \dots, 400, 350)$, and $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_9)$ and $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_9)$ are randomly sampled where the uniform distribution on the set of orthogonal matrices, respectively. We add a sparse outlier matrix $\mathbf{S} = (s_{ij})$ multiplied by $\eta = 1000$, a scaling parameter, to the data matrix \mathbf{X} , i.e.,

$$\mathbf{X} = \mathbf{L} + \mathbf{E}, \quad \mathbf{X}^\eta = \mathbf{L} + \eta \mathbf{S} + \mathbf{E}.$$

The outlier matrix \mathbf{S} has non-zero elements (s_{ij}) only in arbitrarily chosen $0.05n = 50$ rows indexed by I_{50} and $0.05p = 25$ columns indexed by J_{25} . Thus,

$$\begin{cases} s_{ij} = l_{ij} & \text{if } i \in I_{50} \text{ and } j \in J_{25}, \\ s_{ij} = 0 & \text{otherwise.} \end{cases}$$

The simulation is repeated 10 times.

Table 1: Result of The Additional Simulation Study with The Data Size of 2000 by 1000 and Rank 9

METHOD	RIGHT ANGLE	LEFT ANGLE	RATIO OF SINGULAR VALUE	TIME(sec)
SpSVD	4.93	6.11	0.99	17.39
SVD	83.47	81.35	61.42	0.02
RPCA	4.31	6.12	0.96	2987.12
ELSVD	84.92	84.98	63.39	6454.69
R2PCP	77.92	85.49	11.01	890.98
COP	4.19	76.33	18.80	0.09

The results are provided in Table 1. Similarly to the findings in Section 5, SpSVD and RPCA demonstrate superior performance concerning the angles and the ratio, as indicated in the second, third, and fourth columns. COP, similar to our proposal (as it normalizes each vector), exhibits robustness in subspace recovery (right singular subspace). However, it's important to note that the algorithm in Rahmani and Atia (2017) is robust PCA algorithms, not robust SVD algorithms. While they succeed in recovering the right singular vectors (the basis vectors of PC subspaces), they do not provide both left and right singular vectors along with their singular values simultaneously.

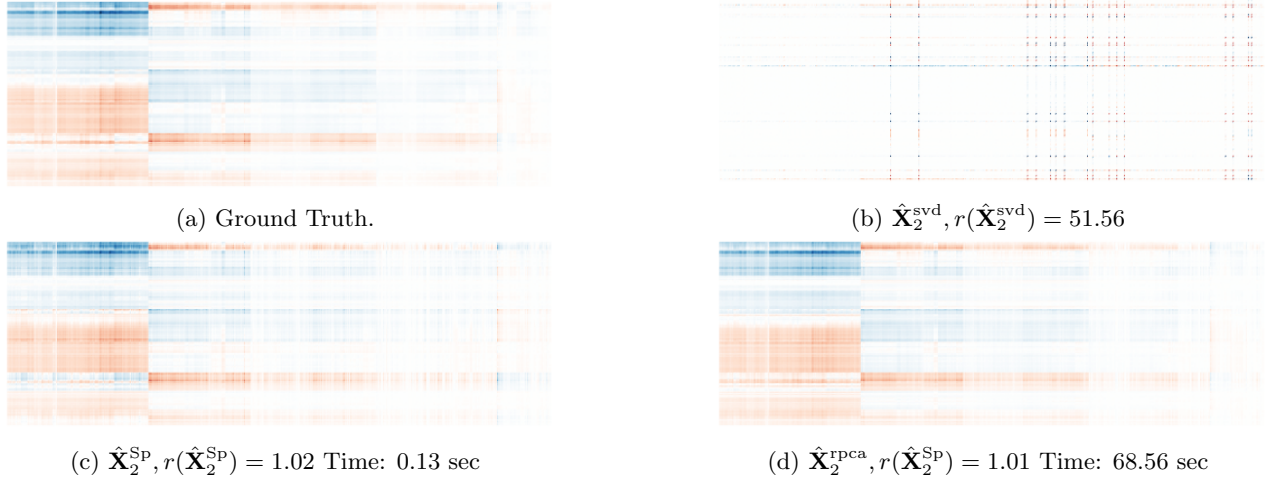


Figure 3: Rank-2 Approximations by SVD, SpSVD and RPCA Using Contaminated Data, Compared with the Ground Truth.

7.3.4 Real Data

The data \mathbf{X} is obtained from the raw data using the following steps. First, each variable in the raw data is transformed using the logarithm base 2 transformation. Then, we apply one-way ANOVA to the centered and normalized raw data with respect to the labels of patients, which have two categories: "good" and "poor", representing the condition of the patients. We select 500 variables from the entire gene expressions based on the p-values obtained from the one-way ANOVA. Subsequently, we obtain the data matrix \mathbf{X} containing these 500 variables and scale it to have zero mean and unit variance.

For the data matrix $\mathbf{X} = (x_{ij})$, we arbitrarily choose approximately $0.1n \approx 16$ rows indexed by I_{16} and $0.1n \approx 16$ columns indexed by J_{16} . We assign outlyingness to the data matrix by amplifying the 16^2 elements corresponding to these rows and columns 1000 times. Thus, the data matrix with contamination, denoted as $\mathbf{X}^\eta = (x_{ij}^\eta)$, is given by

$$\begin{cases} x_{ij}^\eta = 1000x_{ij} & \text{if } i \in I_{16} \text{ and } j \in J_{16}, \\ x_{ij}^\eta = x_{ij} & \text{otherwise.} \end{cases}$$

7.3.5 Lower Bounds of Breakdown Points

In Section 4, we presented the lower bounds n_R and p_R for the breakdown points of the right and left singular spaces obtained from SpSVD, respectively. To determine whether n_R is equal to k for $k = 1, \dots, n$, we need to check the inequality in (8) for all possible submatrices of a row-normalized matrix. However, as the number of rows n increases, this process becomes increasingly challenging, even for small values of k .

We evaluate the lower bounds of the breakdown points using small-sized data. We construct a data matrix $\mathbf{X} = \mathbf{L} + \mathbf{E} \in \mathbb{R}^{n \times p}$ with $n = 200 \times 0.3 = 60$ and $p = 100 \times 0.3 = 30$, where \mathbf{L} is a rank-3 matrix and \mathbf{E} is a matrix consisting of standard normal random noise. The rank-3 matrix \mathbf{L} has singular values $(d_1, d_2, d_3) = (80, 70, 60) \times t$, with $t = 0.3, 0.45, 0.6$, and the corresponding singular vectors are arbitrarily chosen.

Table 2 illustrates the values of the lower bounds $n_R(\mathbf{X})$ and $p_R(\mathbf{X})$ for each $t = 0.3, 0.45, 0.6$ and $R = 3$. When larger singular values are assigned, it implies a larger gap between the two successive singular values of order R and $R + 1$. We observed that the lower bounds n_R and p_R tend to increase as the singular values increase. The lower part of Table 2 also presents similar results for the case where $R = 1$ and the rank-1 matrix \mathbf{L} has a singular value of $60 \times t$ with $t = 0.3, 0.45, 0.6$. Based on the observation where $n = 200 \times 0.3$, $p = 100 \times 0.3$, and

the singular values $(80, 70, 60) \times 0.3$, depicted in the first line of Table 2, we roughly choose the size of the outlier block as $(0.05n, 0.05p)$ in the simulation experiment of Section 5.

Table 2: The Breakdown Points of the Right and Left Singular Spaces Obtained from SpSVD with Various Singular Values and R . Due to the computational burden, we restricted our computations to n_R and p_R up to 6. In cases where the lower bound exceeds or equals 7, we denote it as $7\uparrow$, indicating that an exact value cannot be computed.

The lower bound of breakdown points			
Singular values	R	n_R	p_R
$(80, 70, 60) \times 0.3$	3	3	2
$(80, 70, 60) \times 0.45$	3	5	2
$(80, 70, 60) \times 0.6$	3	6	3
$(60) \times 0.3$	1	6	2
$(60) \times 0.45$	1	$7\uparrow$	4
$(60) \times 0.6$	1	$7\uparrow$	5

7.4 Breakdown points in existing methods

A reviewer suggested to investigate the breakdown points of the existing methods used in the empirical study. The singular vectors from the iterative algorithms of RPCA and R2PCP do not have closed-forms, and it is inherently very challenging to theoretically grasp the breakdown point for these methods. We were not able to find a right technical tool for such purpose. Note that the authors of RPCA and R2PCP did not investigate the breakdown for singular vectors.

It turns out that COP has breakdown points upper-bounded by small numbers. (This is in contrast to our proposal, for which breakdown points are lower-bounded; recall that the higher breakdown point, the more robust a method is.) The COP algorithm, we used in Section 5, has a tuning parameter m , and consists of two steps: Screening out m potential outliers, then applying the vanilla SVD. Below we provide details for rank $R = 1$ SVD approximation.

Let COP_m denote the COP algorithm (that removes m outliers). For data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, let $\mathcal{V}_1^m : \mathbb{R}^{n \times p} \rightarrow \text{Gr}(1, p)$ be given by $\mathcal{V}_1^m(\mathbf{X})$, which is the one-dimensional subspace spanned by the (first) right singular vector of \mathbf{X} obtained by COP_m , and \mathcal{U}_1^m be the left singular vector, obtained by an application of COP_m . In aspects of our breakdown notions, the breakdown points of \mathcal{V}_1^m and \mathcal{U}_1^m for any $\mathbf{X} \in \mathbb{R}^{n \times p}$ are given as

$$\begin{aligned} \text{bp}_{\text{row}}(\mathcal{V}_1^m; \mathbf{X}) &\leq m + 1, \text{bp}_{\text{row}}(\mathcal{U}_1^m; \mathbf{X}) \leq 2, \\ \text{bp}_{\text{col}}(\mathcal{V}_1^m; \mathbf{X}) &\leq 2, \text{bp}_{\text{col}}(\mathcal{U}_1^m; \mathbf{X}) = 1, \\ (k + 1, 2) &\succeq (i, j) \text{ for some } (i, j) \in \text{BP}(\mathcal{V}_1^m; \mathbf{X}), \\ (2, 1) &\succeq (i, j) \text{ for some } (i, j) \in \text{BP}(\mathcal{U}_1^m; \mathbf{X}). \end{aligned}$$

Here, $(2, 1) \succeq (i, j)$ for some $(i, j) \in \text{BP}(\mathcal{U}_1^m(\mathbf{X}))$ implies that block-wise breakdown occurs by changing two elements in one column of \mathbf{X} . Our numerical experiments reflect these theoretical findings: COP_m was shown to be robust and accurate in recovering the right singular subspace (\mathcal{V}_1^m), but failed to recover the left singular subspace (\mathcal{U}_1^m). Note that in these studies we have set m to be the true number of outliers.