

On the Adversarial Robustness of Subspace Learning

Fuwei Li , Lifeng Lai , and Shuguang Cui 

Abstract—In this paper, we study the adversarial robustness of subspace learning problems. Different from the assumptions made in existing works on robust subspace learning where data samples are contaminated by gross sparse outliers or small dense noises, we consider a more powerful adversary who can first observe the data matrix and then intentionally modify the whole data matrix. We first characterize the optimal rank-one attack strategy that maximizes the subspace distance between the subspace learned from the original data matrix and that learned from the modified data matrix. We then generalize the study to the scenario without the rank constraint and characterize the corresponding optimal attack strategy. Besides, our analysis shows that the optimal strategies depend on the singular values of the original data matrix and the adversary's energy budget. Finally, we provide numerical experiments and practical applications to demonstrate the efficiency of the attack strategies.

Index Terms—Subspace learning, principal component analysis, adversarial robustness, non-convex optimization.

I. INTRODUCTION

SUBSPACE learning has a wide range of applications, such as surveillance video analysis, recommendation systems, anomaly detection, etc [2]–[9]. Among a large variety of subspace learning algorithms, principal component analysis (PCA) is one of the most widely used ones. In this paper, we will use PCA as the subspace learning algorithm. PCA computes a small number of principal components, which are orthogonal to each other and represent the majority of the variability of the data samples, and treats the span of these principal components as the desired low-dimensional subspace. Furthermore, many works have proposed robust PCA that can mitigate the impact of certain percentages of outliers and small dense random noise [10]–[13].

In this paper, we investigate the adversarial robustness of subspace learning algorithms. Particularly, we examine the robustness of subspace learning algorithms against not only

random noise or unintentional corrupted data as considered in existing works, but also malicious data produced by powerful adversaries who can modify the whole data set. Our study is motivated by the fact that subspace learning and many other machine learning algorithms are increasingly being used in safety critical and security related applications, such as autonomous vehicle system [14], voice recognition [15], medical image processing [16], etc. In these applications, there might exist powerful adversaries who can modify the data with the goal of maneuvering the machine learning algorithms to make the wrong decision or leave a backdoor in the system [17]. To ensure the security and safety of these systems, it is important to understand the impact of these adversarial attacks on the performance of machine learning algorithms.

In our problem, given the original data matrix, we learn a low-dimensional subspace via PCA. However, there is an adversary who can observe the whole data matrix and then carefully design a modification matrix to change the original data. The goal of the adversary is to modify the original data so as to maximize the subspace distance between the subspace learned from the original data and that learned from the modified data. In this paper, we use Asimov distance [18], defined as the largest principal angle between two subspaces, to measure the subspace distance. Asimov distance has a close relationship with the chordal 2-norm distance and the Finsler distance, which are used in the analysis of optimization on manifolds [19], [20]. It is also related to the gap distance, which is used in the control theory to describe the stability and robustness of a system [21]–[23]. Additionally, it is closely connected to the projection 2-norm that is widely used in various applications [18], [24], [25]. The projection 2-norm provides a way to measure the discrepancy of the projections of a vector on two distinct subspaces. It is useful in the robustness analysis of the principal component regression (PCR), as one is actually projecting the response value vector onto the selected feature subspace in PCR. We will provide an example to illustrate it in Section V using real data. As the Asimov distance depends on the modification matrix in a complex manner, to characterize the optimal attack strategy that maximizes the Asimov distance, we need to solve a complicated non-convex optimization problem.

Towards this goal, we first solve the optimization problem with an additional rank-one constraint on the modification matrix. We note that a rank-one modification is already powerful enough to capture many common modifications such as changing one data sample, inserting one adversarial data point, deleting one feature, etc. Furthermore, the techniques and insights obtained from this special case are useful for the general case without the rank-one constraint. In the rank-one attack case,

Manuscript received August 10, 2019; revised December 31, 2019; accepted February 8, 2020. Date of publication March 4, 2020; date of current version March 11, 2020. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Yue M. Lu. This work was supported in part by National Science Foundation under Grants CCF-1717943, CNS-1824553, and CCF-1908258. This paper was presented in part at IEEE International Conference on Acoustics, Speech, and Signal Processing, Brighton, U.K., May 2019 [1]. (Corresponding author: Shuguang Cui.)

Fuwei Li and Lifeng Lai are with the Department of Electrical and Computer Engineering, University of California, Davis, CA 95616 USA (e-mail: fli@ucdavis.edu; lflai@ucdavis.edu).

Shuguang Cui is with the Shenzhen Research Institute of Big Data and Future Network of Intelligence Institute (FNii), the Chinese University of Hong Kong, Shenzhen 518172, China, and was with the Department of Electrical and Computer Engineering, University of California, Davis, CA 95616 USA (e-mail: robert.cui@gmail.com).

Digital Object Identifier 10.1109/TSP.2020.2974676

we study two different scenarios depending on whether the dimension of the selected subspace is equal to the rank of the data matrix or not. Our study reveals that the optimal attack strategy depends on the energy budget and the singular values of the data matrix. Specifically, in the scenario where the dimension of the selected subspace is the same as the rank of the data matrix, we show that the optimal rank-one strategy depends solely on the energy budget and the smallest singular value of the data matrix. In the scenario where the dimension of the selected subspace is less than the rank of the original data matrix, the optimal strategy depends not only on the energy budget but also on the k th and $(k + 1)$ th singular values, where k is the dimension of the selected subspace.

Relying on the insights gained from the rank-one case, we then extend our study to the more general case where no rank constraint is imposed. Compared with the case with the rank-one constraint, the attacker now has higher degree of freedom to modify the data, which makes the characterization of the optimal attack strategy significantly more challenging. To solve this optimization problem, we first prove that, under the basis of the principal components of the original data matrix, the optimal attack matrix has only few non-zero entries at particular locations. This result greatly reduces the complexity of our problem. With the help of this result, we then simplify our problem to an optimization problem with the objective function being ratio of two quadratic functions. To solve this non-convex problem, we further convert our optimization problem to a feasibility problem and find the closed-form solution to this problem. Our result shows that the optimal strategy depends on the energy budget and the k th and $(k + 1)$ th singular values of the data matrix. Our analysis shows that, compared with the optimal rank-one strategy, this strategy leads to a larger subspace distance.

Our study is related to the recent works on adversarial machine learning. For example, Jagielski *et al.* study how to change the data to manipulate the result of the regression learning system [26]. Bayraktar *et al.* investigate the optimal modification strategy to maximize the inference errors in a multivariate estimation system [27]. In an interesting related work [28], Pimentel-Alarcón *et al.* study how to design an adversarial data sample and add it to the data matrix in order to maximize the Asimov distance between the subspace estimated by PCA from the contaminated data matrix and that from the original data matrix. [28] focuses on the case where the original data matrix is low-rank and the dimension of the selected subspace is equal to the rank of the data matrix. By contrast, we consider a more powerful adversarial setting, where the data matrix is not constrained to being low-rank, the dimension of the selected subspace does not necessarily equal the rank of the data matrix, and the adversary can modify the whole data matrix instead of only adding one data sample.

The remainder of the paper is organized as follows. In Section II, we describe the precise problem formulation. In Section III, we investigate the optimal rank-one attack strategy. We generalize our results to the case without the rank constraint in Section IV. In Section V, we provide numerical experiments with both synthesized data and real data to illustrate results

obtained in this paper. Finally, we offer concluding remarks in Section VI.

II. PROBLEM FORMULATION

In this section, we introduce the problem formulation. Given a data matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$ with each $\mathbf{x}_i \in \mathbb{R}^d$, our goal is to learn a low-dimension subspace via PCA. In the data matrix \mathbf{X} , we assume that all the preprocessing steps (such as data centering and standardization) have been done. In this paper, we consider an adversarial setup in which an adversary will first observe \mathbf{X} and then carefully design a modification (attack) matrix $\Delta\mathbf{X}$ to change \mathbf{X} to $\hat{\mathbf{X}} = \mathbf{X} + \Delta\mathbf{X}$. We denote function $g_k(\cdot)$ as the PCA operation that computes the k leading principal components. Furthermore, let $\mathbb{X} = \text{span}(g_k(\mathbf{X}))$ be a k -dimensional subspace learned from \mathbf{X} and $\hat{\mathbb{X}} = \text{span}(g_k(\hat{\mathbf{X}}))$ be a k -dimensional subspace learned from the modified data matrix $\hat{\mathbf{X}}$. The goal of the adversary is to design the modification matrix $\Delta\mathbf{X}$ so as to make the distance between \mathbb{X} and $\hat{\mathbb{X}}$ as large as possible. To measure such a distance, we use the largest principal angle between \mathbb{X} and $\hat{\mathbb{X}}$ as defined below [18].

Definition 1: Let \mathbb{X} and $\hat{\mathbb{X}}$ be two k -dimensional subspaces in \mathbb{R}^d . The principal angles $\{\theta_i\}_{i=1}^k$ are defined recursively:

$$\begin{aligned} \cos(\theta_i) &= \max_{\mathbf{u}_i \in \mathbb{X}, \mathbf{v}_i \in \hat{\mathbb{X}}} \mathbf{u}_i^\top \mathbf{v}_i \\ \text{s.t.} \quad & \|\mathbf{u}_i\| = \|\mathbf{v}_i\| = 1, \\ & \mathbf{u}_j^\top \mathbf{u}_i = \mathbf{v}_j^\top \mathbf{v}_i = 0, \forall j = 1, 2, \dots, i-1. \end{aligned}$$

In this paper, we will use $\|\cdot\|$ to denote the ℓ_2 norm and $\theta(g_k(\mathbf{X}), g_k(\hat{\mathbf{X}}))$ or simply θ to denote the Asimov distance between the subspace \mathbb{X} estimated from \mathbf{X} and the subspace $\hat{\mathbb{X}}$ estimated from $\hat{\mathbf{X}}$. Given an orthonormal basis $\mathbf{U}_{\mathbb{X}}$ of \mathbb{X} and an orthonormal basis $\mathbf{U}_{\hat{\mathbb{X}}}$ of $\hat{\mathbb{X}}$, $\{\cos(\theta_1), \dots, \cos(\theta_k)\}$ are the singular values of $\mathbf{U}_{\mathbb{X}}^\top \mathbf{U}_{\hat{\mathbb{X}}}$ [18]. Hence, the Asimov distance is determined by the smallest singular value of $\mathbf{U}_{\mathbb{X}}^\top \mathbf{U}_{\hat{\mathbb{X}}}$. It is easy to see that, if no constraint is imposed on $\Delta\mathbf{X}$, $\hat{\mathbb{X}}$ can be arbitrary and θ can be easily made to be $\pi/2$. Therefore, we impose an energy constraint on $\Delta\mathbf{X}$. In particular, we assume that the energy of $\Delta\mathbf{X}$ is less than or equal to η . In this paper, we use the Frobenius norm $\|\Delta\mathbf{X}\|_F$ to measure the energy. Hence, the goal of this attacker is to solve the following optimization problem:

$$\begin{aligned} \max_{\Delta\mathbf{X} \in \mathbb{R}^{d \times n}} & : \theta(g_k(\mathbf{X}), g_k(\hat{\mathbf{X}})) \\ \text{s.t.} \quad & \hat{\mathbf{X}} = \mathbf{X} + \Delta\mathbf{X}, \\ & \|\Delta\mathbf{X}\|_F \leq \eta. \end{aligned} \quad (1)$$

Even though (1) is a complicated non-convex optimization problem, we will fully characterize the optimal solution to (1) for any given η . This characterization will enable us to investigate the impact of this optimal attack with respect to the energy budget η .

Note that we consider a very powerful adversary model who has access to the whole dataset and can modify all data points. For security analysis, it is desirable to consider the worst case

scenario with powerful adversary. Furthermore, our analysis provides a universal upper bound on the maximum subspace distance incurred by any bounded energy perturbation.

III. OPTIMAL RANK-ONE ADVERSARIAL STRATEGY

In this section, we will solve (1) for the special case where the modification matrix $\Delta\mathbf{X}$ is limited to being rank-one. The techniques and insights obtained from this special case will be useful for the general case considered in Section IV.

With this additional rank-one constraint, $\Delta\mathbf{X}$ can be written as $\mathbf{a}\mathbf{b}^\top$ for some $\mathbf{a} \in \mathbb{R}^d$ and $\mathbf{b} \in \mathbb{R}^n$, and the optimization problem (1) becomes

$$\begin{aligned} \max_{\mathbf{a} \in \mathbb{R}^d, \mathbf{b} \in \mathbb{R}^n} : & \theta(g_k(\mathbf{X}), g_k(\hat{\mathbf{X}})) \\ \text{s.t.} & \hat{\mathbf{X}} = \mathbf{X} + \Delta\mathbf{X}, \\ & \Delta\mathbf{X} = \mathbf{a}\mathbf{b}^\top, \\ & \|\Delta\mathbf{X}\|_F \leq \eta. \end{aligned} \quad (2)$$

It is easy to see that, for any feasible solution $(\tilde{\mathbf{a}}, \tilde{\mathbf{b}})$ with $\|\tilde{\mathbf{b}}\| \neq 1$, we can construct another feasible solution $(\|\tilde{\mathbf{b}}\|\tilde{\mathbf{a}}, \tilde{\mathbf{b}}/\|\tilde{\mathbf{b}}\|)$ that gives the same objective function value. Hence, without loss of optimality, we will fix the norm of \mathbf{b} to be 1 throughout this section.

Based on the value of k , i.e., the dimension of the subspace we select, we will first present the solution to the case when $k = \text{rank}(\mathbf{X})$, and then generalize the result to the case when $k < \text{rank}(\mathbf{X})$.

A. Case With $k = \text{Rank}(\mathbf{X})$

In this subsection, we consider the case when the dimension of the subspace selected is equal to the rank of the data matrix. In this case, the span of \mathbf{X} equals the span of $g_k(\mathbf{X})$. Furthermore, we divide this case into two scenarios where the data matrix is full-rank and the data matrix is low-rank.

1) *Full-Rank Case*: In the full column rank case, $\text{rank}(\mathbf{X}) = n$, where $n \leq d$. This case arises when the number of samples is limited, for example, at the beginning of online PCA. In this case, the span of $\hat{\mathbf{X}}$ is equal to the span of $g_k(\hat{\mathbf{X}})$, and hence we can write $\theta(g_k(\mathbf{X}), g_k(\hat{\mathbf{X}}))$ as $\theta(\mathbf{X}, \hat{\mathbf{X}})$. In the following, we first find the expression of $\theta(\mathbf{X}, \hat{\mathbf{X}})$ for any given $\hat{\mathbf{X}} = \mathbf{X} + \mathbf{a}\mathbf{b}^\top$. Using this expression, we then characterize the optimal attack matrix $\Delta\mathbf{X}$.

Suppose the compact SVD of \mathbf{X} is $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top = \mathbf{U}\mathbf{W}$, where $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_n)$. One set of orthonormal bases for the column space of \mathbf{X} is \mathbf{U} . We can also use SVD to find a set of orthonormal bases $\tilde{\mathbf{U}}$ of span($\hat{\mathbf{X}}$).

Since $\hat{\mathbf{X}} = \mathbf{X} + \mathbf{a}\mathbf{b}^\top$, $\tilde{\mathbf{U}}$ can be directly expressed as a function of \mathbf{U} [29]:

$$\tilde{\mathbf{U}} = \mathbf{U} + (\alpha\mathbf{U}\mathbf{w} + \beta\mathbf{s})\mathbf{w}^\top,$$

where

$$\begin{aligned} \mathbf{a}_{u^\perp} &= (\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{a}, & \mathbf{s} &= \mathbf{a}_{u^\perp}/\|\mathbf{a}_{u^\perp}\|, \\ \tilde{\mathbf{w}} &= -\mathbf{W}^{-\top}\mathbf{b}, & \mathbf{w} &= \tilde{\mathbf{w}}/\|\tilde{\mathbf{w}}\|, \\ \omega &= (1 - \mathbf{a}^\top\mathbf{U}\tilde{\mathbf{w}})/\|\mathbf{a}_{u^\perp}\|, & \mathbf{g} &= [\tilde{\mathbf{w}}, \omega]^\top, \end{aligned}$$

$$\alpha = |\omega|/\|\mathbf{g}\| - 1, \quad \beta = -\text{sign}(\omega)\|\tilde{\mathbf{w}}\|/\|\mathbf{g}\|,$$

and $\mathbf{W}^{-\top} = (\mathbf{W}^{-1})^\top$. Hence, we have $\mathbf{U}^\top\tilde{\mathbf{U}} = \mathbf{U}^\top(\mathbf{U} + (\alpha\mathbf{U}\mathbf{w} + \beta\mathbf{s})\mathbf{w}^\top) = \mathbf{I} + \alpha\mathbf{w}\mathbf{w}^\top$. The singular values of $\mathbf{I} + \alpha\mathbf{w}\mathbf{w}^\top$ are $\{1, 1, \dots, 1 + \alpha\mathbf{w}^\top\mathbf{w}\}$. Since $\mathbf{w}^\top\mathbf{w} = 1$, $1 + \alpha = |\omega|/\|\mathbf{g}\|$, the smallest singular value of $\mathbf{U}^\top\tilde{\mathbf{U}}$ is $\cos(\theta) = |\omega|/\|\mathbf{g}\|$. Our objective is to maximize θ , which is equivalent to minimizing the smallest singular value of $\mathbf{U}^\top\tilde{\mathbf{U}}$. Hence, the optimization problem (2) is simplified as

$$\begin{aligned} \min_{\mathbf{a}, \mathbf{b}} : & |\omega|/\|\mathbf{g}\| \\ \text{s.t.} & \|\mathbf{a}\mathbf{b}^\top\|_F = \|\mathbf{a}\|\|\mathbf{b}\| \leq \eta, \end{aligned}$$

where we use the identity $\|\mathbf{a}\|\|\mathbf{b}\| = \|\mathbf{a} \cdot \mathbf{b}^\top\|_F$. Expanding the objective function, we have

$$\frac{|\omega|}{\|\mathbf{g}\|} = \frac{|1 + \mathbf{a}_u^\top\mathbf{W}^{-\top}\mathbf{b}|}{\|[\|\mathbf{a}_{u^\perp}\|\mathbf{W}^{-\top}\mathbf{b}, 1 + \mathbf{a}_u^\top\mathbf{W}^{-\top}\mathbf{b}]\|}, \quad (3)$$

where $\mathbf{a}_u = \mathbf{U}^\top\mathbf{a}$.

Since $\mathbf{W} = \Sigma\mathbf{V}^\top$, we have $\mathbf{W}^{-\top}\mathbf{b} = \Sigma^{-1}\mathbf{V}^\top\mathbf{b}$. As \mathbf{V} is a unitary matrix, changing the coordinate $\mathbf{b} \leftarrow \mathbf{V}^\top\mathbf{b}$ does not result in the change of the constraint. The value $\mathbf{a}_u^\top\mathbf{W}^{-\top}\mathbf{b}$ in the original coordinate is the same as $\mathbf{a}_u^\top\Sigma^{-1}\mathbf{b}$ in the new coordinate. In the following, we will use this new coordinate system and the cost function in (3) can be written as

$$\frac{|\omega|}{\|\mathbf{g}\|} = \frac{|1 + \mathbf{a}_u^\top\Sigma^{-1}\mathbf{b}|}{\|[\|\mathbf{a}_{u^\perp}\|\Sigma^{-1}\mathbf{b}, 1 + \mathbf{a}_u^\top\Sigma^{-1}\mathbf{b}]\|}. \quad (4)$$

The objective function (4) is zero if and only if the numerator is zero. Using the matrix norm inequality [30], we have

$$\begin{aligned} |\mathbf{a}_u^\top\Sigma^{-1}\mathbf{b}| &\leq \|\mathbf{a}_u\|\|\mathbf{b}\|\|\Sigma^{-1}\|_2 = \frac{1}{\sigma_n}\|\mathbf{a}_u\|\|\mathbf{b}\| \\ &\stackrel{(a)}{\leq} \frac{1}{\sigma_n}\|\mathbf{a}\|\|\mathbf{b}\| = \frac{1}{\sigma_n}\|\mathbf{a}\mathbf{b}^\top\|_F \stackrel{(b)}{\leq} \frac{\eta}{\sigma_n}, \end{aligned}$$

where $\|\Sigma^{-1}\|_2$ is the induced 2-norm of matrix Σ^{-1} , in (a) we use $\|\mathbf{a}_u\| \leq \|\mathbf{a}\|$, and (b) is due to the energy constraint. From the inequalities, we conclude that when $\eta < \sigma_n$, we can not make the numerator to be zero. We now consider two different cases depending on whether we can make the numerator to be zero or not.

Case 1: When $\eta > \sigma_n$, if we set

$$\mathbf{a}_u = [0, 0, \dots, -\sigma_n]^\top, \quad \mathbf{b} = [0, 0, \dots, 1]^\top,$$

and any $\|\mathbf{a}_{u^\perp}\|^2 = \hat{a}^2$ with $0 < \hat{a}^2 < \eta^2 - \sigma_n^2$, the numerator will be zero. Since $\mathbf{a} = \mathbf{U}\mathbf{a}_u + (\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{a}_{u^\perp}$, the attacker can make the Asimov distance to be $\pi/2$ by setting:

$$\mathbf{a} = -\sigma_n\mathbf{u}_n + \hat{a}\mathbf{u}_q, \quad \mathbf{b} = \mathbf{v}_n, \quad (5)$$

where \mathbf{u}_q is any vector orthogonal to the column space of \mathbf{X} and $0 < \hat{a}^2 < \eta^2 - \sigma_n^2$.

Case 2: When $\eta \leq \sigma_n$, the value of $1 + \mathbf{a}_u^\top\Sigma^{-1}\mathbf{b}$ can not reach zero. In this case, it is easy to check that minimizing (4) is equivalent to maximizing

$$\frac{\|\mathbf{a}_{u^\perp}\|^2\|\Sigma^{-1}\mathbf{b}\|^2}{(1 + \mathbf{a}_u^\top\Sigma^{-1}\mathbf{b})^2}. \quad (6)$$

As $\|\mathbf{b}\| = 1$, $\|\Sigma^{-1}\mathbf{b}\|^2$ is maximized when $\mathbf{b} = [0, 0, \dots, 1]^\top$. Furthermore, for any fixed norm of \mathbf{a}_u , $(1 + \mathbf{a}_u^\top \Sigma^{-1}\mathbf{b})^2$ is minimized when $\mathbf{a}_u = [0, 0, \dots, -\|\mathbf{a}_u\|]^\top$, $\mathbf{b} = [0, 0, \dots, 1]^\top$. Hence, for fixed norms of \mathbf{a}_u , \mathbf{a}_{u^\perp} , the objective function (6) is maximized when

$$\mathbf{a}_u = [0, 0, \dots, -\|\mathbf{a}_u\|]^\top, \quad \mathbf{b} = [0, 0, \dots, 1]^\top. \quad (7)$$

Let $c = \|\mathbf{a}_{u^\perp}\|$, $h = \|\mathbf{a}_u\|$. Using the optimal form of \mathbf{a}_u and \mathbf{b} in (7), the objective function (6) can be simplified to

$$\begin{aligned} \max_{c,h} : & \frac{c^2/\sigma_n^2}{(1-h/\sigma_n)^2} \\ \text{s.t.} & (c^2 + h^2) \leq \eta^2, \end{aligned} \quad (8)$$

It is easy to check that the objective function is maximized when $c^2 + h^2 = \eta^2$. Hence, we have $c^2 = \eta^2 - h^2$. Inserting this value of c into the objective function and setting the derivative with respect to h to be 0, we get a unique solution $h = \eta^2/\sigma_n$. At this value of h , the second derivative is $\frac{-2\sigma_n^2}{(\sigma_n^2 - \eta^2)^3}$, which is negative. It indicates that $h = \eta^2/\sigma_n$ is indeed the maximum point. Hence, $c = \pm\eta\sqrt{1 - \eta^2/\sigma_n^2}$. This implies that the optimal solution to problem (2) for Case 2 is

$$\mathbf{a} = -\eta^2/\sigma_n \mathbf{u}_n \pm \eta\sqrt{1 - \eta^2/\sigma_n^2} \mathbf{u}_q, \quad \mathbf{b} = \mathbf{v}_n.$$

Summarizing the discussion above, we have the following proposition regarding the optimal value of problem (2) in the full-rank case.

Proposition 1: In the full rank case, the optimal value of (2) is

$$\theta^* = \begin{cases} \pi/2, & \text{if } \eta > \sigma_n \\ \arcsin(\eta/\sigma_n), & \text{if } \eta \leq \sigma_n \end{cases}.$$

2) *Low-Rank Case:* We now consider the case where \mathbf{X} is not full rank. Let $k < \min(d, n)$ be the rank of \mathbf{X} . In this subsection, with a slight abuse of notation, we write the full SVD of \mathbf{X} as $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top$. The optimal attack matrix could be found by solving

$$\begin{aligned} \max_{\mathbf{a} \in \mathbb{R}^d, \mathbf{b} \in \mathbb{R}^n} : & \theta(\mathbf{X}, g_k(\hat{\mathbf{X}})) \\ \text{s.t.} & \hat{\mathbf{X}} = \mathbf{X} + \mathbf{a}\mathbf{b}^\top, \\ & \|\mathbf{a}\|\|\mathbf{b}\| \leq \eta. \end{aligned} \quad (9)$$

We can further simplify this optimization problem as

$$\begin{aligned} \max_{\mathbf{a} \in \mathbb{R}^{k+1}, \mathbf{b} \in \mathbb{R}^{k+1}} : & \theta(\tilde{\Sigma}, g_k(\mathbf{Y})) \\ \text{s.t.} & \mathbf{Y} = \tilde{\Sigma} + \mathbf{a}\mathbf{b}^\top, \\ & \|\mathbf{a}\|\|\mathbf{b}\| \leq \eta, \end{aligned} \quad (10)$$

where $\tilde{\Sigma} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_k, 0)$ and $\{\sigma_1, \sigma_2, \dots, \sigma_k\}$ are singular values of \mathbf{X} . Detailed proof of the equivalence between (9) and (10) can be found in Appendix A. Here, we describe the main idea of the proof. The main step of the simplification is to left multiply the unitary matrix \mathbf{U}^\top and right multiply the unitary matrix \mathbf{V} on both $\hat{\mathbf{X}}$ and $\tilde{\mathbf{X}}$. Note that

multiplying a unitary matrix does not change the column space and its singular values. In addition, a rank-one modification can only add at most one principal component orthogonal to its original column subspace. Hence, by changing the coordinates, \mathbf{a} and \mathbf{b} are $k+1$ dimensional vectors.

To solve problem (10), we divide it into two cases based on the value of the energy budget.

Case 1: When $\eta > \sigma_k$, it is simple to verify that the solution $\mathbf{a} = [0, 0, \dots, \eta]^\top$, $\mathbf{b} = [0, 0, \dots, 1]^\top$ leads to the maximal Asimov distance, which is $\pi/2$.

Case 2: When $\eta \leq \sigma_k$, the following theorem characterizes the form of optimal \mathbf{a} and \mathbf{b} .

Theorem 1: There exists an optimal solution to problem (10) in the following form

$$\mathbf{a} = [0, \dots, 0, a_k, a_{k+1}]^\top, \quad \mathbf{b} = [0, 0, \dots, 0, 1, 0]^\top, \quad (11)$$

with $a_k^2 + a_{k+1}^2 = \eta^2$.

Proof: Please see Appendix B. ■

In the following, we will find the optimal values of a_k and a_{k+1} . Since $\|\mathbf{a}\|^2 = \eta^2$ and \mathbf{a} is in the form of (11), we can write $\mathbf{a} = \eta[0, 0, \dots, \cos(\alpha), \sin(\alpha)]^\top$, where $\alpha \in [0, 2\pi)$. To compute the k leading principal components of \mathbf{Y} , we can perform the eigenvalue decomposition of $\mathbf{Y}\mathbf{Y}^\top$,

$$\mathbf{Y}\mathbf{Y}^\top = \begin{bmatrix} \Lambda_{k-1}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{c}\mathbf{c}^\top \end{bmatrix},$$

where $\mathbf{c} = [\sigma_k + \eta \cos \alpha, \eta \sin(\alpha)]^\top$, $\Lambda_{k-1} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_{k-1})$. Suppose the compact SVD of $\mathbf{Y}\mathbf{Y}^\top$ is $\mathbf{Y}\mathbf{Y}^\top = \hat{\mathbf{U}}\hat{\Sigma}\hat{\mathbf{V}}^\top$, where

$$\hat{\mathbf{U}} = \begin{bmatrix} \mathbf{I}_{k-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{z} \end{bmatrix},$$

and $\mathbf{z} \in \mathbb{R}^2$ is the eigenvector of $\mathbf{c}\mathbf{c}^\top$ corresponding to its nonzero eigenvalue. Since one orthonormal basis of $\text{span}(\tilde{\Sigma})$ is $[\mathbf{I}_k, \mathbf{0}]^\top$, the Asimov distance is determined by the singular values of

$$\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix}^\top \cdot \begin{bmatrix} \mathbf{I}_{k-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{z} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{k-1} & \mathbf{0} \\ \mathbf{0} & z_1 \end{bmatrix}.$$

Hence, the Asimov distance is $\arccos(|z_1|)$. Since \mathbf{c} is the eigenvector of $\mathbf{c}\mathbf{c}^\top$ corresponding to its nonzero eigenvalue, we have $|z_1| = \frac{|\mathbf{c}_1|}{\|\mathbf{c}\|}$. Our objective function is reduced to

$$\min_{\alpha \in [0, 2\pi)} : \frac{|\sigma_k + \eta \cos(\alpha)|}{\|[\sigma_k + \eta \cos(\alpha), \eta \sin(\alpha)]\|}. \quad (12)$$

It is simple to show that the optimal solution to (12) is

$$\alpha^* = \arccos(-\eta/\sigma_k) \quad (13)$$

or

$$\alpha^* = 2\pi - \arccos(-\eta/\sigma_k). \quad (14)$$

Substitute the optimal solution of α^* in (13) or (14) into the objective of problem (12), we have $\sin(\theta^*) = \eta/\sigma_k$. Hence, the

optimal solution to problem (10) is

$$\mathbf{a} = \left[0, 0, \dots, -\eta^2/\sigma_k, \pm\eta\sqrt{1-\eta^2/\sigma_k^2} \right]^\top,$$

$$\mathbf{b} = [0, 0, \dots, 0, 1, 0]^\top,$$

which indicates that the optimal solution to problem (9) is

$$\mathbf{a} = -\eta^2/\sigma_k \mathbf{u}_k \pm \eta\sqrt{1-\eta^2/\sigma_k^2} \mathbf{u}_q, \quad \mathbf{b} = \mathbf{v}_k,$$

where \mathbf{u}_q is any vector orthogonal to the column space of \mathbf{X} . The corresponding optimal subspace distance is $\theta^* = \arcsin(\eta/\sigma_k)$. In summary, we have

Proposition 2: The optimal Asimov distance in the low-rank case is

$$\theta^* = \begin{cases} \pi/2, & \text{if } \eta > \sigma_k \\ \arcsin(\eta/\sigma_k), & \text{if } \eta \leq \sigma_k \end{cases}. \quad (15)$$

The result is similar to the full column rank case characterized in Proposition 1.

B. Case With $k < \text{Rank}(\mathbf{X})$

In this section, we consider the more practical but much more challenging case with $k < \text{rank}(\mathbf{X})$.

Given the data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$, without loss of generality, we assume $d \leq n$ and $\text{rank}(\mathbf{X}) = d$. Assume the full SVD of \mathbf{X} is $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, where $\mathbf{U} \in \mathbb{R}^{d \times d}$, $\mathbf{\Sigma} \in \mathbb{R}^{d \times n}$, $\mathbf{V} \in \mathbb{R}^{n \times n}$, and the singular values of \mathbf{X} are $\{\sigma_1, \sigma_2, \dots, \sigma_k, \dots, \sigma_d\}$. Recall that we denote $g_k(\cdot)$ as the PCA operation that computes the k leading principal components. In this scenario, as the original data matrix is not low-rank, we will perform PCA both on the original data matrix and on the modified data matrix. Hence, the optimal rank-one modification matrix can be found by solving the following optimization problem

$$\begin{aligned} \max_{\mathbf{a} \in \mathbb{R}^d, \mathbf{b} \in \mathbb{R}^n} &: \theta(g_k(\mathbf{X}), g_k(\hat{\mathbf{X}})) \\ \text{s.t.} & \quad \hat{\mathbf{X}} = \mathbf{X} + \mathbf{a}\mathbf{b}^\top, \\ & \quad \|\mathbf{a}\mathbf{b}^\top\|_F \leq \eta. \end{aligned} \quad (16)$$

By diagonalizing the data matrix and using similar arguments in Appendix A, (16) can be further simplified as

$$\begin{aligned} \max_{\mathbf{a} \in \mathbb{R}^d, \mathbf{b} \in \mathbb{R}^n} &: \theta(g_k(\mathbf{\Sigma}), g_k(\mathbf{Y})) \\ \text{s.t.} & \quad \mathbf{Y} = \mathbf{\Sigma} + \mathbf{a}\mathbf{b}^\top, \\ & \quad \|\mathbf{a}\mathbf{b}^\top\|_F \leq \eta, \end{aligned} \quad (17)$$

where $g_k(\mathbf{\Sigma}) = [\mathbf{I}_k, \mathbf{0}]^\top \in \mathbb{R}^{d \times k}$. Here we also perform variable change $\mathbf{a} \leftarrow \mathbf{U}^\top \mathbf{a}$ and $\mathbf{b} \leftarrow \mathbf{V}^\top \mathbf{b}$. To solve this optimization problem, we divide it into two cases depending on the energy budget and the difference between σ_k and σ_{k+1} .

Case 1: When $\eta \geq \sigma_k - \sigma_{k+1}$, we have one simple solution $\mathbf{a} = [0, 0, \dots, 0, \eta, 0, \dots, 0]^\top$, where η is in the $(k+1)$ th coordinate, and $\mathbf{b} = [0, 0, \dots, 0, 1, 0, \dots, 0]^\top$, where element 1 is in the $(k+1)$ th coordinate. Clearly, this setting of \mathbf{a} and \mathbf{b} leads to the maximal subspace distance, which is $\pi/2$.

Case 2: When $\eta < \sigma_k - \sigma_{k+1}$, the following theorem gives the form of the optimal solution.

Theorem 2: The optimal solution to problem (17) should be in the form of

$$\mathbf{a} = [0, 0, \dots, a_k, a_{k+1}, 0, \dots, 0]^\top, \quad (18)$$

$$\mathbf{b} = [0, 0, \dots, b_k, b_{k+1}, 0, \dots, 0]^\top, \quad (19)$$

where $a_k^2 + a_{k+1}^2 = \eta^2$ and $b_k^2 + b_{k+1}^2 = 1$.

Proof: Please see Appendix C for details. \blacksquare

As the optimal solution of \mathbf{a} and \mathbf{b} is in the form of (18) and (19), we can parametrize \mathbf{a} and \mathbf{b} with parameters α and β using $\mathbf{a} = \eta[0, 0, \dots, \cos(\alpha), \sin(\alpha), 0, \dots, 0]^\top$ and $\mathbf{b} = [0, 0, \dots, \cos(\beta), \sin(\beta), 0, \dots, 0]^\top$ respectively.

As a result, the modified data matrix \mathbf{Y} can be written as

$$\mathbf{Y} = \begin{bmatrix} \mathbf{\Sigma}_1 & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_2 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{\Sigma}_3 & \mathbf{0} \end{bmatrix},$$

where $\mathbf{\Sigma}_1 = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_{k-1})$, $\mathbf{\Sigma}_3 = \text{diag}(\sigma_{k+2}, \dots, \sigma_d)$, and

$$\mathbf{\Sigma}_2 = \begin{bmatrix} \sigma_k + \eta \cos(\alpha) \cos(\beta) & \eta \cos(\alpha) \sin(\beta) \\ \eta \sin(\alpha) \cos(\beta) & \sigma_{k+1} + \eta \sin(\alpha) \sin(\beta) \end{bmatrix}. \quad (20)$$

Since \mathbf{Y} has the pseudo block diagonal form, the singular values and principal components of \mathbf{Y} are determined by the SVD of $\mathbf{\Sigma}_1$, $\mathbf{\Sigma}_2$, and $\mathbf{\Sigma}_3$. For notation convenience, we denote $\mathbf{\Sigma}_2 = \mathbf{D} + \eta \bar{\mathbf{a}} \bar{\mathbf{b}}^\top$, where $\mathbf{D} = \text{diag}(\sigma_k, \sigma_{k+1})$, $\bar{\mathbf{a}} = [\cos \alpha, \sin \alpha]^\top$, and $\bar{\mathbf{b}} = [\cos \beta, \sin \beta]^\top$. Let ξ_1 and ξ_2 be the two singular values of $\mathbf{\Sigma}_2$ and denote their corresponding left singular vectors as

$$\mathbf{W} = [\mathbf{w}_1 \ \mathbf{w}_2] = \begin{bmatrix} \cos \varphi & -\sin \varphi \\ \sin \varphi & \cos \varphi \end{bmatrix}. \quad (21)$$

The following lemma characterizes the form of the k -dimensional subspace learned by PCA from \mathbf{Y} .

Lemma 1:

$$g_k(\mathbf{Y}) = \begin{bmatrix} \mathbf{I}_{k-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{w}_1 \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

Proof: According to the perturbation theory [31], the singular values of $\mathbf{\Sigma}_2$ must satisfy

$$\xi_2 < \sigma_k, \quad \xi_1 > \sigma_{k+1}.$$

It indicates that $\xi_1 > \sigma_k$, $\xi_2 > \sigma_k$ and $\xi_1 < \sigma_{k+1}$, $\xi_2 < \sigma_{k+1}$ will not happen. Hence, we will select the eigenvector corresponding to singular value ξ_1 as one of the leading k principal components, which completes the proof. \blacksquare

Since one set of orthonormal bases for $g_k(\mathbf{\Sigma})$ is $[\mathbf{I}_k, \mathbf{0}]^\top$, the subspace distance $\theta(g_k(\mathbf{\Sigma}), g_k(\mathbf{Y}))$ is determined by the singular values of

$$\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix}^\top \cdot \begin{bmatrix} \mathbf{I}_{k-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{w}_1 \\ \mathbf{0} & \mathbf{0} \end{bmatrix} = \text{diag}(1, 1, \dots, \cos \varphi).$$

Hence, the subspace distance is $\arccos(|\cos \varphi|)$ and our optimization problem can be equivalently formulated as

$$\min_{\alpha \in [0, 2\pi], \beta \in [0, 2\pi]} |\cos \varphi|. \quad (22)$$

Let $\mathbf{Z} = \Sigma_2 \Sigma_2^\top$, we can compute \mathbf{W} through eigenvalue decomposition of \mathbf{Z} . According to the equality $\Sigma_2 \Sigma_2^\top = \mathbf{W} \cdot \text{diag}(\xi_1^2, \xi_2^2) \cdot \mathbf{W}^\top$, we have

$$\begin{aligned} \mathbf{Z} &= \begin{bmatrix} Z_{1,1} & Z_{1,2} \\ Z_{2,1} & Z_{2,2} \end{bmatrix} \\ &= \begin{bmatrix} \xi_1^2 \cos^2 \varphi + \xi_2^2 \sin^2 \varphi & (\xi_1^2 - \xi_2^2) \cos \varphi \sin \varphi \\ (\xi_1^2 - \xi_2^2) \cos \varphi \sin \varphi & \xi_1^2 \sin^2 \varphi + \xi_2^2 \cos^2 \varphi \end{bmatrix}. \end{aligned}$$

From this equation, we obtain

$$\begin{cases} \cos(2\varphi)(\xi_1^2 - \xi_2^2) = Z_{1,1} - Z_{2,2} \\ \sin(2\varphi)(\xi_1^2 - \xi_2^2) = Z_{1,2} + Z_{2,1} \end{cases}.$$

Then we can compute φ through

$$\varphi = 0.5 \text{atan2}(a_y, a_x), \quad (23)$$

where $\text{atan2}(\cdot, \cdot)$ is the four-quadrant inverse tangent function, $a_x = Z_{1,1} - Z_{2,2}$, and $a_y = Z_{1,2} + Z_{2,1}$. In our case, the specific expressions of a_x and a_y are

$$\begin{cases} a_x = \sigma_k^2 - \sigma_{k+1}^2 + 2\sigma_k \eta \cos(\alpha) \cos(\beta) \\ \quad - 2\sigma_{k+1} \eta \sin(\alpha) \sin(\beta) + \eta^2 \cos(2\alpha), \\ a_y = 2\eta (\sigma_k \sin(\alpha) \cos(\beta) + \sigma_{k+1} \cos(\alpha) \sin(\beta) \\ \quad + \eta \cos(\alpha) \sin(\alpha)). \end{cases} \quad (24)$$

Let us write a_x and a_y as a function of α and β : $a_x = a_x(\alpha, \beta)$ and $a_y = a_y(\alpha, \beta)$. To further restrict the domains of α and β , we analyze the properties of the angle φ in (23) as a function of α and β . First, we have $a_x(\alpha, \beta) = a_x(\pi + \alpha, \pi + \beta)$ and $a_y(\alpha, \beta) = a_y(\pi + \alpha, \pi + \beta)$. So $\varphi(\alpha, \beta) = \varphi(\pi + \alpha, \pi + \beta)$. This property indicates that we only need to consider the function value in the domain $\alpha \in [0, \pi], \beta \in [-\pi, \pi]$. Second, $a_x(\alpha, \beta) = a_x(\pi - \alpha, \pi - \beta)$ and $a_y(\alpha, \beta) = -a_y(\pi - \alpha, \pi - \beta)$, and then we have $\varphi(\alpha, \beta) = -\varphi(\pi - \alpha, \pi - \beta)$. Since $\cos(\varphi)$ is an even function, we only need to consider the function with domain $\alpha \in [0, \pi/2], \beta \in [-\pi, \pi]$. Note that Σ_2 is in the form of (20), the variance in the direction of \mathbf{e}_k is $v_k = \cos(\alpha)^2 + \sigma_k^2 + 2\cos(\alpha)\cos(\beta)$, and the variance in the direction of \mathbf{e}_{k+1} is $v_{k+1} = \sin(\alpha)^2 + \sigma_{k+1}^2 + 2\sin(\alpha)\sin(\beta)$. To maximize the subspace distance, we should make v_k small and make v_{k+1} large. Apparently, the sign of $\cos(\alpha)\cos(\beta)$ should be negative and the sign of $\sin(\alpha)\sin(\beta)$ should be positive. Hence, the optimal α and β should satisfy $\alpha \in [0, \pi/2]$ and $\beta \in [\pi/2, \pi]$. As a result, the optimization problem (22) can be written as

$$\min_{\alpha \in [0, \pi/2], \beta \in [\pi/2, \pi]} |\cos(\varphi(\alpha, \beta))|. \quad (25)$$

The following theorem characterizes the optimal solution to problem (25).

Theorem 3: The optimal solution to problem (25) is

$$\begin{cases} \alpha^* = \arccos\left(\sqrt{\frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 - \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)}}\right), \\ \beta^* = \arccos\left(-\sqrt{\frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 + \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)}}\right), \end{cases} \quad (26)$$

where $H = \sigma_k^4 + \sigma_{k+1}^4 + \eta^4 - 2\sigma_k^2\sigma_{k+1}^2 - 2\sigma_k^2\eta^2 - 2\sigma_{k+1}^2\eta^2$.

Proof: Please see Appendix D. ■

Accordingly, the optimal solution to problem (16) is

$$\mathbf{a}^* = \eta \cos(\alpha^*) \mathbf{u}_k + \eta \sin(\alpha^*) \mathbf{u}_{k+1}, \quad (27)$$

$$\mathbf{b}^* = \cos(\beta^*) \mathbf{v}_k + \sin(\beta^*) \mathbf{v}_{k+1}. \quad (28)$$

Furthermore, the optimal subspace distance θ^* can be computed according to (24) and (23). Moreover, according to the properties of the function $\varphi(\alpha, \beta)$ we have discussed before, there are other three optimal solutions

$$(-\alpha^*, -\beta^*), \quad (\pi - \alpha^*, \pi - \beta^*), \quad (\alpha^* - \pi, \beta^* - \pi),$$

which lead to the same optimal objective value.

IV. OPTIMAL ADVERSARIAL STRATEGY WITHOUT THE RANK CONSTRAINT

Using the insights gained from Section III, we now characterize the optimal attack strategy in the general case without the rank-one constraint by solving (1). We will directly consider the general case with $k \leq \text{rank}(\mathbf{X})$.

Following the similar transformation from (9) to (10), we can simplify the optimization problem (1) as

$$\begin{aligned} \max_{\mathbf{B} \in \mathbb{R}^{d \times n}} &: \theta(g_k(\Sigma), g_k(\mathbf{Y})) \\ \text{s.t.} & \quad \mathbf{Y} = \Sigma + \mathbf{B}, \\ & \quad \|\mathbf{B}\|_F \leq \eta, \end{aligned} \quad (29)$$

where without loss of generality we assume $d \leq n$, the full SVD of the data matrix is $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top$, the singular values of the data matrix are $\{\sigma_1, \sigma_2, \dots, \sigma_d\}$, and $\mathbf{B} = \mathbf{U}^\top \Delta \mathbf{X} \mathbf{V}$. To identify the optimal modification matrix \mathbf{B} in problem (29), we divide it into two cases.

Case 1: When $\eta \geq \frac{\sigma_k - \sigma_{k+1}}{\sqrt{2}}$, by setting $b_{k,k} = -\eta/\sqrt{2}$, $b_{k+1,k+1} = \eta/\sqrt{2}$, and all other entries of \mathbf{B} to zero, where $b_{i,j}$ is the element in the i th row and j th column of \mathbf{B} , it will lead to the maximal subspace distance, $\pi/2$.

Case 2: When $\eta < \frac{\sigma_k - \sigma_{k+1}}{\sqrt{2}}$, the following theorem states the form of the optimal \mathbf{B} .

Theorem 4: The optimal \mathbf{B} to problem (29) has only four possible non-zero entries: $b_{k,k}$, $b_{k,k+1}$, $b_{k+1,k}$ and $b_{k+1,k+1}$.

Proof: Please see Appendix E. ■

This characterization reduces the complexity of problem (29). Using this optimal form of \mathbf{B} and following similar steps leading to (23), we can write the subspace distance as

$$\theta = 0.5 |\text{atan2}(b_y, b_x)|, \quad (30)$$

where

$$\begin{aligned} b_y &= 2((b_{k,k} + \sigma_k)b_{k+1,k} + (b_{k+1,k+1} + \sigma_{k+1})b_{k,k+1}), \\ b_x &= (b_{k,k} + \sigma_k)^2 + b_{k,k+1}^2 - (b_{k+1,k+1} + \sigma_{k+1})^2 - b_{k+1,k}^2. \end{aligned}$$

It is easy to see that we can change the sign of b_y by changing the signs of $b_{k,k+1}$ and $b_{k+1,k}$. We also have $b_x > 0$, as

$$\begin{aligned} & \frac{b_x}{\| [b_{k,k} + \sigma_k, b_{k,k+1}] \| + \| [b_{k+1,k+1} + \sigma_{k+1}, b_{k+1,k}] \|} \\ &= \| [b_{k,k} + \sigma_k, b_{k,k+1}] \| - \| [b_{k+1,k+1} + \sigma_{k+1}, b_{k+1,k}] \| \\ &\geq \sigma_k - \sigma_{k+1} - \| [b_{k,k}, b_{k,k+1}] \| - \| [b_{k+1,k}, b_{k+1,k+1}] \| \\ &\geq \sigma_k - \sigma_{k+1} - \sqrt{2}\eta > 0. \end{aligned}$$

Using these two facts and the fact that $\text{atan2}(b_y, b_x)$ is an odd function of b_y when $b_x > 0$, we know that maximizing θ in (30) is equivalent to maximizing b_y/b_x . Hence, our optimization problem can be written as

$$\begin{aligned} \max_{\mathbf{u}} : & \frac{\mathbf{u}^\top \mathbf{A}_1 \mathbf{u}}{\mathbf{u}^\top \mathbf{A}_2 \mathbf{u}} \\ \text{s.t.} & \quad \| \mathbf{u} - \boldsymbol{\sigma} \|^2 \leq \eta^2, \end{aligned} \quad (31)$$

where $\mathbf{u} \triangleq \bar{\mathbf{b}} + \boldsymbol{\sigma}$ with $\bar{\mathbf{b}} = [b_{k,k}, b_{k+1,k}, b_{k,k+1}, b_{k+1,k+1}]^\top$ and $\boldsymbol{\sigma} = [\sigma_k, 0, 0, \sigma_{k+1}]^\top$,

$$\mathbf{A}_1 = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad \text{and} \quad \mathbf{A}_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}.$$

The objective function is the ratio of two quadratic functions. It is a non-convex problem in general. In the following, we transform this problem into a feasibility problem and obtain the closed-form solution analytically.

Let λ denote the value of the objective function in (31). We can rewrite the optimization problem (31) as

$$\begin{aligned} \max_{\lambda, \mathbf{u}} : & \quad \lambda \\ \text{s.t.} & \quad \frac{\mathbf{u}^\top \mathbf{A}_1 \mathbf{u}}{\mathbf{u}^\top \mathbf{A}_2 \mathbf{u}} = \lambda, \\ & \quad \| \mathbf{u} - \boldsymbol{\sigma} \|^2 \leq \eta^2. \end{aligned} \quad (32)$$

The first constraint can be written as $\mathbf{u}^\top (\mathbf{A}_1 - \lambda \mathbf{A}_2) \mathbf{u} = 0$, where

$$\begin{bmatrix} \mathbf{Q} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q} \end{bmatrix} \triangleq \mathbf{A}_1 - \lambda \mathbf{A}_2 = \begin{bmatrix} -\lambda & 1 & 0 & 0 \\ 1 & \lambda & 0 & 0 \\ 0 & 0 & -\lambda & 1 \\ 0 & 0 & 1 & \lambda \end{bmatrix}.$$

To further simplify the constraint, we perform eigenvalue decomposition on $\mathbf{Q} = \mathbf{P} \boldsymbol{\Lambda} \mathbf{P}^\top$, where $\boldsymbol{\Lambda} = \text{diag}(\sqrt{\lambda^2 + 1}, -\sqrt{\lambda^2 + 1})$ and

$$\mathbf{P} = t \begin{bmatrix} 1 & -(\sqrt{\lambda^2 + 1} + \lambda) \\ \sqrt{\lambda^2 + 1} + \lambda & 1 \end{bmatrix}, \quad (33)$$

with $t = 1/\sqrt{(\sqrt{\lambda^2 + 1} + \lambda)^2 + 1}$.

We further perform variable change $\mathbf{v} \triangleq \text{diag}(\mathbf{P}^\top, \mathbf{P}^\top) \mathbf{u}$. Thus, the constraint (32) is equivalent to $\mathbf{v}^\top \boldsymbol{\Lambda} \mathbf{v} = 0$, which indicates $v_1^2 + v_3^2 = v_2^2 + v_4^2$. With this, the optimization problem is simplified as

$$\max_{\lambda, \mathbf{v}} : \quad \lambda \quad (34)$$

$$\text{s.t.} \quad v_1^2 + v_3^2 = v_2^2 + v_4^2, \quad (35)$$

$$\| \mathbf{v} - \bar{\boldsymbol{\sigma}} \|^2 \leq \eta^2, \quad (36)$$

where $\bar{\boldsymbol{\sigma}} = \text{diag}(\mathbf{P}^\top, \mathbf{P}^\top) \boldsymbol{\sigma} = [p_{1,1}\sigma_k, p_{1,2}\sigma_k, p_{2,1}\sigma_{k+1}, p_{2,2}\sigma_{k+1}]^\top$. Note that $p_{1,2} = -p_{2,1}$ and $p_{2,2} = p_{1,1}$, we have $\bar{\boldsymbol{\sigma}} = [p_{1,1}\sigma_k, -p_{2,1}\sigma_k, p_{2,1}\sigma_{k+1}, p_{1,1}\sigma_{k+1}]^\top$.

Now, problem (34) can be solved by checking the feasibility of (35) and (36) given a particular λ . Given λ , the feasibility of problem (34) is equivalent to the feasibility of

$$\min_{v_1^2+v_3^2=v_2^2+v_4^2} \| \mathbf{v} - \bar{\boldsymbol{\sigma}} \|^2 \leq \eta^2. \quad (37)$$

Note that $\bar{\boldsymbol{\sigma}}$ depends on λ , we denote the left hand side of inequality (37) as $f(\mathbf{v}, \lambda) = \| \mathbf{v} - \bar{\boldsymbol{\sigma}} \|^2$ and parametrize \mathbf{v} as

$$v_1 = r \cos(\alpha), v_2 = r \cos(\beta), v_3 = r \sin(\alpha), v_4 = r \sin(\beta). \quad (38)$$

It is easy to verify that the minimum point of $f(\mathbf{v}, \lambda)$ in terms of \mathbf{v} is obtained at the following stationary point

$$\begin{cases} r = \frac{1}{2} \left(\sqrt{p_{1,1}^2 \sigma_k^2 + p_{2,1}^2 \sigma_{k+1}^2} + \sqrt{p_{2,1}^2 \sigma_k^2 + p_{1,1}^2 \sigma_{k+1}^2} \right), \\ \cos(\alpha) = p_{1,1} \sigma_k / \sqrt{p_{1,1}^2 \sigma_k^2 + p_{2,1}^2 \sigma_{k+1}^2}, \\ \sin(\alpha) = p_{2,1} \sigma_{k+1} / \sqrt{p_{1,1}^2 \sigma_k^2 + p_{2,1}^2 \sigma_{k+1}^2}, \\ \cos(\beta) = -p_{2,1} \sigma_k / \sqrt{p_{1,1}^2 \sigma_{k+1}^2 + p_{2,1}^2 \sigma_k^2}, \\ \sin(\beta) = p_{1,1} \sigma_{k+1} / \sqrt{p_{1,1}^2 \sigma_{k+1}^2 + p_{2,1}^2 \sigma_k^2}. \end{cases} \quad (39)$$

Plug the optimal r, α, β of (39) into $f(\mathbf{v}, \lambda)$, and we have

$$\begin{aligned} f(\lambda) &\triangleq \min_{v_1^2+v_3^2=v_2^2+v_4^2} f(\mathbf{v}, \lambda) \\ &= (\sigma_k^2 + \sigma_{k+1}^2)/2 \\ &\quad - \sqrt{p_{1,1}^2 \sigma_k^2 + p_{2,1}^2 \sigma_{k+1}^2} \sqrt{p_{2,1}^2 \sigma_k^2 + p_{1,1}^2 \sigma_{k+1}^2}. \end{aligned}$$

According to inequality (37), inequality $f(\lambda) \leq \eta^2$ now is equivalent to

$$\begin{aligned} & \sqrt{p_{1,1}^2 \sigma_k^2 + p_{2,1}^2 \sigma_{k+1}^2} \sqrt{p_{2,1}^2 \sigma_k^2 + p_{1,1}^2 \sigma_{k+1}^2} \\ & \geq (\sigma_k^2 + \sigma_{k+1}^2)/2 - \eta^2. \end{aligned} \quad (40)$$

Denote the right hand of the above inequality as $c \triangleq (\sigma_k^2 + \sigma_{k+1}^2)/2 - \eta^2$. Since $\eta < (\sigma_k - \sigma_{k+1})/\sqrt{2}$, we have $c > \sigma_k \sigma_{k+1}$. Furthermore, we notice that $p_{1,1}^2 = 1 - p_{2,1}^2$. Plug it into inequality (40), and we have

$$p_{2,1}^4 - p_{2,1}^2 + \frac{c^2 - \sigma_k^2 \sigma_{k+1}^2}{(\sigma_k^2 - \sigma_{k+1}^2)^2} \leq 0. \quad (41)$$

Let

$$w \triangleq \frac{c^2 - \sigma_k^2 \sigma_{k+1}^2}{(\sigma_k^2 - \sigma_{k+1}^2)^2}, \quad (42)$$

and since $\sigma_k \sigma_{k+1} < c \leq (\sigma_k^2 + \sigma_{k+1}^2)/2$, we have $0 < w \leq \frac{(\sigma_k^2 + \sigma_{k+1}^2)^2/4 - \sigma_k^2 \sigma_{k+1}^2}{(\sigma_k^2 - \sigma_{k+1}^2)^2} = 1/4$. Denote the left hand of inequality (41) as $h(p_{2,1})$, and we have

$$\begin{aligned} h_{\min} &= h(1/\sqrt{2}) = -1/4 + w \leq 0, \\ h(1) &= w > 0. \end{aligned}$$

Moreover, since $1/\sqrt{2} < p_{2,1} < 1$, we must have

$$p_{2,1} \leq p_{2,1}^H, \quad (43)$$

where $p_{2,1}^H = \sqrt{(1 + \sqrt{1 - 4w})/2}$ is the largest root of $h(p_{2,1}) = 0$. Plugging the expressions of $p_{2,1}$ and $p_{2,1}^H$ into (43), we can get

$$\frac{\sqrt{\lambda^2 + 1} + \lambda}{\sqrt{(\sqrt{\lambda^2 + 1} + \lambda)^2 + 1}} \leq \sqrt{\frac{1 + \sqrt{1 - 4w}}{2}}.$$

Simplifying this inequality leads to $\lambda \leq \frac{e^2 - 1}{2e}$, where

$$e = \sqrt{\frac{1 + \sqrt{1 - 4w}}{1 - \sqrt{1 - 4w}}}. \quad (44)$$

Thus we can conclude that

$$\lambda_{\max} = \frac{e^2 - 1}{2e}. \quad (45)$$

Accordingly, the optimal subspace distance in (1) is

$$\theta^* = \text{atan}(\lambda_{\max})/2. \quad (46)$$

In summary, given energy budget η , we first compute w according to (42) and compute e according to (44), from which we can get λ_{\max} and θ^* using (45) and (46). Having obtained the optimal λ_{\max} , we can compute \mathbf{P} in (33) and compute \mathbf{v} using (39) and (38), and sequentially compute \mathbf{u} and $\bar{\mathbf{b}}$. Finally, if the optimal solution of problem (29) is \mathbf{B}^* with non-zero entries $\bar{\mathbf{b}}^* = [b_{k,k}^*, b_{k+1,k}^*, b_{k,k+1}^*, b_{k+1,k+1}^*]^\top$, we also have another paired feasible optimal solution with non-zero entries being $[b_{k,k}^*, -b_{k+1,k}^*, -b_{k,k+1}^*, b_{k+1,k+1}^*]^\top$, which leads to the same optimal value. Accordingly, the optimal solution to problem (1) is $\Delta \mathbf{X}^* = \mathbf{U} \mathbf{B}^* \mathbf{V}^\top$.

V. NUMERICAL EXPERIMENTS AND APPLICATIONS

In this section, we provide numerical examples to illustrate the results obtained in this paper. We will also apply the results to principal component regression [32] to illustrate potential applications in practice.

A. Numerical Experiments

In this subsection, we illustrate the results with synthesized data.

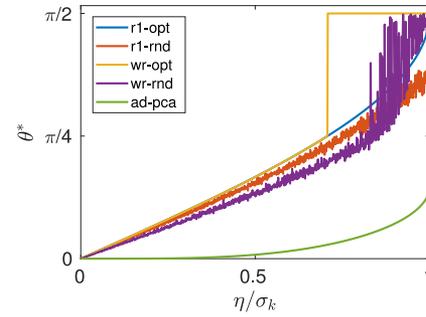


Fig. 1. Subspace distances with different attack strategies on a low-rank data matrix over different energy budgets.

In the first experiment, we employ different attack strategies in a low-rank data matrix. In this simulation, we set $d = 5$, $n = 5$, and $k = 3$. We generate the original data matrix as $\mathbf{X} = \mathbf{A} \mathbf{B}^\top$, where $\mathbf{A} \in \mathbb{R}^{d \times k}$, $\mathbf{B} \in \mathbb{R}^{n \times k}$, and each entry of \mathbf{A} and \mathbf{B} is i.i.d. generated according to a standard normal distribution. First, we conduct our optimal rank-one attack strategy. In this strategy, we use the result from the analysis of the optimal rank-one modification matrix to design \mathbf{a} , \mathbf{b} and add the attack matrix $\Delta \mathbf{X} = \mathbf{a} \mathbf{b}^\top$ to the original data matrix \mathbf{X} . We then perform SVD on $\hat{\mathbf{X}}$ and select the k leading principal components. Finally, we compute the distance between the selected subspace and the original subspace. We also conduct a test using a random rank-one attack strategy, in which we randomly generate \mathbf{a} , \mathbf{b} with each entry of \mathbf{a} , \mathbf{b} being i.i.d. generated according to the standard normal distribution. Then we normalize the energy of $\mathbf{a} \mathbf{b}^\top$ to be η^2 . For each η , we repeatedly generate 100000 pairs of \mathbf{a} and \mathbf{b} and compute their corresponding subspace distances. In addition, we compare it with the strategy where the modification matrix is free of rank constraint. Although our analysis is deliberately designed for general data matrices, we set the $(k + 1)$ th singular value to be zero so that it can be applied to the low-rank data matrix. We design the modification matrix $\Delta \mathbf{X}$ according to our analysis in this paper and calculate the subspace distance between the original subspace and that after modification. Moreover, we conduct another random attack strategy in which we randomly generate the modification matrix without any rank constraint. Each entry of the modification matrix is i.i.d. generated according to a standard normal distribution. After that, we normalize its Frobenius norm equal to η . We repeat this attack 100000 times for each η and record its corresponding subspace distance. Furthermore, we also compare it with the strategy described in [28], which adds one adversarial data sample into the data set.

Fig. 1 demonstrates the subspace distances obtained by the five strategies. In this figure, r1-opt represents the rank-one optimal attack obtained in this paper, r1-rnd represents the maximal subspace distance obtained among the 100000 times random rank-one attacks, wr-opt stands for our optimal attack without the rank constraint, wr-rnd is the maximal subspace distance among the 100000 random attacks without the rank constraint, and ad-pca is the algorithm described in [28]. The x axis is the ratio between η and the smallest singular value of the original data matrix. From the figure, we can see our optimal strategies

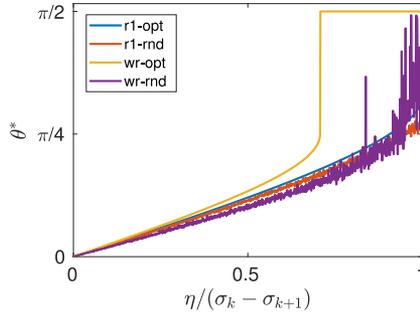


Fig. 2. Subspace distances achieved by using different attack strategies under different energy budgets.

are much better than the ad-pca strategy. It is because our strategies can modify the data matrix, and thus have higher degree of freedom to manipulate the data. The optimal strategies designed in this paper also have a larger subspace distance compared with their corresponding random attack strategies. In the region where $\eta/\sigma_k \in [0, 1/\sqrt{2}]$, both of our two optimal strategies provide the same subspace distances, which can be verified by setting $\sigma_{k+1} = 0$, computing θ^* in equation (46) and comparing it with the value in equation (15). When $\eta/\sigma_k > 1/\sqrt{2}$, the optimal attack without the rank constraint leads to the largest subspace distance, $\pi/2$, which is much larger than the distance obtained by the optimal rank-one attack strategy. That means, without the rank constraint, it indeed provides a larger subspace distance.

In the second numerical experiment, we test these strategies except the ad-pca in the general data matrix in which the data matrix is not low-rank. In this experiment, we set $d = 5$, $n = 5$, and $k = 3$. We randomly generate the data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$ with each entries i.i.d generated according to a standard normal distribution. We also design the optimal rank-one attack matrix and the optimal modification matrix without the rank constraint according to the analysis provided in this paper. In addition, we do random attacks 100000 times using the randomly generated modification matrix with the rank-one constraint and without the rank constraint respectively.

Fig. 2 shows the subspace distances obtained through different strategies over different energy budgets. In this figure, the x axis is the ratio between η and $\sigma_k - \sigma_{k+1}$. For the two random attack strategies, we demonstrate the maximal subspace distances achieved by the 100000 times random attacks. As the figure shows, both of the two random strategies have smaller subspace distances compared with their perspective optimal strategies. Different from the low-rank case, the strategy without the rank constraint provides larger subspace distances consistently over all the energy budgets.

B. Applications

In this subsection, we use real data to illustrate the results obtained in this paper.

In particular, we illustrate the impact of adversarial attack on PCR, which is widely used in statistical learning especially when collinearity exists in the data. Ordinary regression will increase

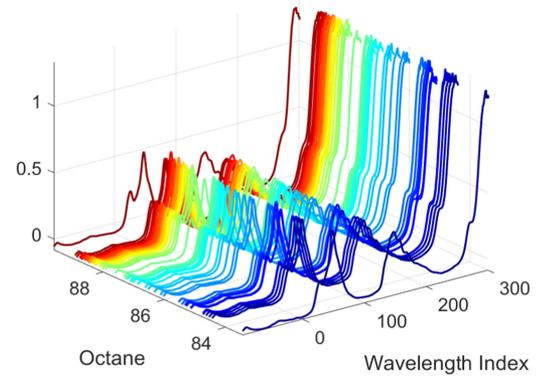


Fig. 3. The spectral intensities of the gasoline data set.

the standard error of the coefficients when there are high correlations or even collinearities between features. This happens particularly when the number of features is much larger than the number of data samples. PCR deals with this issue by performing PCA on the feature matrix and only selecting the leading k principal components as the predictors, and thus dramatically decreases the number of predictors. The regression process of PCR can be seen as projecting the response values onto the subspace spanned by the leading k principal components. So, the accuracy of the subspace will significantly influence the regression results. Appendix F provides an example on how the change of the subspace will influence the result of PCR. More details of PCR can be found in [32].

In this experiment, our task is to use the gasoline spectral intensity to predict its octane rating. We use the gasoline spectral data set [33], which comprises spectral intensities of 60 samples of gasoline at 401 wavelengths, and their octane ratings. Fig. 3 shows the spectral intensities of the data set. This figure indicates that the correlation of intensity among different wavelengths is very high. To complete the regression task, we can use PCR.

In this experiment, we randomly select 80 percent of the data as the training set and the remaining 20 percent as the test set. We choose 4 principal components as our predictors and perform regression based on these principal components. We also record the r-squared values both in the training phase and the test phase. The r-squared value is defined as $r^2 = 1 - \frac{\|\mathbf{y} - \hat{\mathbf{y}}\|^2}{\|\mathbf{y} - \bar{\mathbf{y}}\|^2}$, where r^2 is the r-squared value, \mathbf{y} is the response values, $\hat{\mathbf{y}}$ is the predicted values, $\|\mathbf{y} - \bar{\mathbf{y}}\|^2$ represents the total variance of the response values, and $\bar{\mathbf{y}} = \text{mean}(\mathbf{y}) \cdot \mathbf{1}$ stands for the mean vector of the response values. R-squared value measures how well the model fits the data and larger r-squared value indicates better regression. Firstly, we perform regular PCR without attack and let na-train and na-test denote the r-squared values of the training and test respectively. We then attack the feature matrix using the optimal rank-one strategy proposed in this paper with different energies and denote r1-train and r1-test as its r-squared values in the training and test processes. Finally, we also carry out the optimal attack without the rank constraint and denote wr-train, wr-test as the r-squared values in the training and test procedures.

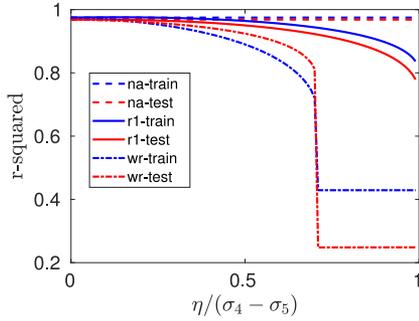


Fig. 4. R-squared values with different attack strategies over different energy budgets.

Fig. 4 illustrates the r-squared values with different attack strategies under different energy budgets. As shown in this figure, with the increasing of the energy budget, r-squared values of training and test decrease for both attack strategies. This figure also indicates that the strategy with no rank constraint is more efficient than the rank-one strategy considering its smaller r-squared values. Furthermore, the r-squared value of the strategy without the rank constraint has a tremendous drop at the point $\eta/(\sigma_4 - \sigma_5) = 1/\sqrt{2}$, which is consistent with our analysis that beyond this particular point the maximal subspace distance is $\pi/2$.

VI. CONCLUSION

In this paper, we have investigated the adversarial robustness of PCA problem. We have characterized the optimal rank-one adversarial modification strategy and the optimal strategy without the rank constraint to modify the data. Our analysis has showed that both of the two strategies depend on the singular values of the data matrix and the adversary's energy budget. We have also performed numerical simulations and investigated the impact of this attack on PCR. Both the numerical experiments and the PCR application illustrate that adversarial attacks degrade the performance of subspace learning significantly. In the future, it is of interest to investigate the defense strategy to mitigate the effects of this attack.

APPENDIX A

PROOF OF THE EQUIVALENCE BETWEEN PROBLEM (9) AND PROBLEM (10)

Before giving the proof, we first examine the unitary invariant property of the Asimov distance, which is helpful in our subsequent proof.

Proposition 3: Let \mathbf{P} and \mathbf{T} be unitary matrices, and then for the Asimov distance function $\theta(\cdot, \cdot)$, we have

$$\theta(\mathbf{X}_1, g_k(\mathbf{X}_2)) = \theta(\mathbf{P}\mathbf{X}_1\mathbf{T}^\top, g_k(\mathbf{P}\mathbf{X}_2\mathbf{T}^\top)).$$

Proof: First, we show $\theta(\mathbf{X}_1, \mathbf{X}_2) = \theta(\mathbf{P}\mathbf{X}_1\mathbf{T}^\top, \mathbf{P}\mathbf{X}_2\mathbf{T}^\top)$. Suppose the thin QR decompositions of \mathbf{X}_1 and \mathbf{X}_2 are $\mathbf{X}_1 = \mathbf{Q}_1\mathbf{R}_1$, $\mathbf{X}_2 = \mathbf{Q}_2\mathbf{R}_2$, and then the subspace distance between the two subspaces spanned by the columns of \mathbf{X}_1 and \mathbf{X}_2 is determined by the singular values of $\mathbf{Q}_1^\top\mathbf{Q}_2$. Since $(\mathbf{P}\mathbf{Q}_1)^\top(\mathbf{P}\mathbf{Q}_2) = \mathbf{Q}_1^\top\mathbf{Q}_2$ and right multiplying a unitary matrix does not change

the singular values and the column subspace of a matrix, we have $\theta(\mathbf{X}_1, \mathbf{X}_2) = \theta(\mathbf{P}\mathbf{X}_1\mathbf{T}^\top, \mathbf{P}\mathbf{X}_2\mathbf{T}^\top)$.

Second, suppose the full SVD of \mathbf{X}_2 is $\mathbf{X}_2 = \mathbf{U}_2\mathbf{\Sigma}_2\mathbf{V}_2^\top$, where $\mathbf{U}_2 = [\mathbf{u}_{21}, \mathbf{u}_{22}, \dots, \mathbf{u}_{2d}]$. Then

$$\mathbf{P}g_k(\mathbf{X}_2) = \mathbf{P}[\mathbf{u}_{21}, \mathbf{u}_{22}, \dots, \mathbf{u}_{2k}] = g_k(\mathbf{P}\mathbf{X}_2),$$

which can be verified by checking that $\mathbf{P}\mathbf{U}_2\mathbf{\Sigma}_2\mathbf{V}_2^\top$ is a valid SVD of $\mathbf{P}\mathbf{X}_2$. It completes the proof. \blacksquare

With the help of this proposition, let $\mathbf{P} = \mathbf{U}^\top$, $\mathbf{T} = \mathbf{V}^\top$, right multiply \mathbf{P} and left multiply \mathbf{T}^\top on both \mathbf{X} and $\hat{\mathbf{X}}$, and we can simplify problem (9) as the following

$$\begin{aligned} \max_{\mathbf{a} \in \mathbb{R}^d, \mathbf{b} \in \mathbb{R}^n} & : \theta(\mathbf{\Sigma}, g_k(\tilde{\mathbf{Y}})) \\ \text{s.t.} & \quad \tilde{\mathbf{Y}} = \mathbf{\Sigma} + \mathbf{a}\mathbf{b}^\top, \\ & \quad \|\mathbf{a}\| \|\mathbf{b}\| \leq \eta, \end{aligned} \quad (47)$$

where we assume $n > d$, $\mathbf{\Sigma} = [\text{diag}(\sigma_1, \sigma_2, \dots, \sigma_k, \mathbf{0}), \mathbf{0}] \in \mathbb{R}^{d \times n}$. Also, from problem (9) to problem (47), we do variable change $\mathbf{a} \leftarrow \mathbf{U}^\top \mathbf{a}$, $\mathbf{b} \leftarrow \mathbf{V}^\top \mathbf{b}$.

To further simplify this optimization problem, we split \mathbf{a} and \mathbf{b} into $\mathbf{a} = [\mathbf{a}_1^\top, \mathbf{a}_2^\top]^\top$, $\mathbf{b} = [\mathbf{b}_1^\top, \mathbf{b}_2^\top]^\top$, where $\mathbf{a}_1 \in \mathbb{R}^k$, $\mathbf{a}_2 \in \mathbb{R}^{d-k}$, $\mathbf{b}_1 \in \mathbb{R}^k$, and $\mathbf{b}_2 \in \mathbb{R}^{n-k}$. In addition, utilizing the Householder transformation [30], we construct an orthogonal matrix

$$\mathbf{M}_1 = \begin{bmatrix} \mathbf{I}_k & \mathbf{0} \\ \mathbf{0} & \mathbf{H}_1 \end{bmatrix}, \quad (48)$$

where

$$\begin{aligned} \mathbf{M}_1^\top \mathbf{M}_1 &= \mathbf{I}, & \mathbf{H}_1 &= \mathbf{I} - 2 \frac{\mathbf{u}\mathbf{u}^\top}{\|\mathbf{u}\|^2}, \\ \mathbf{u} &= \mathbf{a}_2 - s_1 \|\mathbf{a}_2\| \cdot \mathbf{e}_1, & \mathbf{e}_1 &= [1, 0, \dots, 0]^\top \in \mathbb{R}^{d-k}, \\ \mathbf{H}_1^\top \mathbf{a}_2 &= s_1 \|\mathbf{a}_2\| \cdot \mathbf{e}_1, & s_1 &= \pm 1. \end{aligned}$$

Similarly, we can construct another Householder transformation matrix \mathbf{H}_2 for \mathbf{b}_2 and the corresponding orthogonal matrix $\mathbf{M}_2 = \text{diag}(\mathbf{I}_k, \mathbf{H}_2)$. Left multiplying \mathbf{M}_1^\top and right multiplying \mathbf{M}_2 on $\tilde{\mathbf{Y}}$, we have

$$\mathbf{M}_1^\top \tilde{\mathbf{Y}} \mathbf{M}_2 = \begin{bmatrix} \tilde{\mathbf{\Sigma}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{a}_1 \\ s_1 \|\mathbf{a}_2\| \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{b}_1^\top & s_2 \|\mathbf{b}_2\| & \mathbf{0} \end{bmatrix},$$

where $s_2 = \pm 1$.

Let $\mathbf{a} \leftarrow [\mathbf{a}_1^\top, s_1 \|\mathbf{a}_2\|]^\top$ and $\mathbf{b} \leftarrow [\mathbf{b}_1^\top, s_2 \|\mathbf{b}_2\|]^\top$. Utilizing Proposition 3, it is clear that problem (10) and problem (47) are equivalent.

APPENDIX B

PROOF OF THEOREM 1

The proof follows similar steps to those in [28]. In problem (10), $\tilde{\mathbf{\Sigma}}$ is a diagonal matrix with diagonal elements $\{\sigma_1, \sigma_2, \dots, \sigma_k, 0\}$. The subspace spanned by $g_k(\mathbf{Y})$ is a k -dimensional subspace in \mathbb{R}^{k+1} . We denote this subspace as \mathbb{Q} , denote \mathbb{P} as the subspace spanned by $\tilde{\mathbf{\Sigma}}$ and further denote their intersection as $\mathbb{T} = \mathbb{P} \cap \mathbb{Q}$. Note that \mathbb{P} is not

equal to \mathbb{Q} (otherwise the Asimov distance will be zero), so we have $\dim(\mathbb{P} \cup \mathbb{Q}) = k + 1$. Since $\dim(\mathbb{P}) + \dim(\mathbb{Q}) - \dim(\mathbb{T}) = \dim(\mathbb{P} \cup \mathbb{Q})$, we have $\dim(\mathbb{T}) = k - 1$. Let \mathbf{T} be an orthonormal basis of \mathbb{T} . Let $[\mathbf{T}, \mathbf{p}]$ be an orthonormal basis of \mathbb{P} and let $[\mathbf{T}, \mathbf{q}]$ be an orthonormal basis of \mathbb{Q} . By the definition of Asimov distance, the subspace distance between \mathbb{P} and \mathbb{Q} is the angle between \mathbf{p} and \mathbf{q} .

Firstly, it is easy to see that $a_{k+1} \neq 0$. Otherwise, \mathbb{Q} will be equal to \mathbb{P} , which means that their Asimov distance is zero.

Secondly, it is easy to see $\mathbf{q} \in \text{span}[\mathbf{T}, \mathbf{p}, \mathbf{e}_{k+1}]$, where \mathbf{e}_{k+1} is an ordinary basis vector that only has element 1 in the $(k+1)$ th coordinate. Since \mathbf{T} is orthogonal to \mathbf{q} , we have $\mathbf{q} \in \text{span}[\mathbf{p}, \mathbf{e}_{k+1}]$. It is easy to see that the larger variance in the direction of \mathbf{p} is, the closer \mathbf{p} and \mathbf{q} will be. Then we should select \mathbf{p} as the direction with the smallest variance in \mathbf{X} . Since we are assuming that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k$, \mathbf{p} should be \mathbf{e}_k .

Thirdly, for a fixed direction of \mathbf{a} , let $\hat{\mathbf{a}}$ be the projection of \mathbf{a} onto $\text{span}[\mathbf{e}_k, \mathbf{e}_{k+1}]$. Clearly, \mathbf{q} will be closer to \mathbf{a} as $\hat{\mathbf{a}}$ grows. As a result, the angle between \mathbf{q} and \mathbf{p} will be larger. This also implies that the length of \mathbf{a} should be maximized: $\|\mathbf{a}\| = \eta$. Hence, the Asimov distance is maximized when $\mathbf{a} = \hat{\mathbf{a}}$ and $\|\mathbf{a}\| = \eta$, implying that \mathbf{a} only has nonzero elements in its k th and $k+1$ th coordinates.

Finally, for a fixed \mathbf{a} in the form of (11), the projected variance of \mathbf{Y} on the direction of \mathbf{e}_k is $v_1 = \sum_{i \neq k} (a_k b_i)^2 + (a_k b_k + \sigma_k)^2 = a_k^2 + \sigma_k^2 + 2a_k b_k \sigma_k$ and the projected variance of \mathbf{Y} on the direction of \mathbf{e}_{k+1} is $v_2 = \sum_i (a_{k+1} b_i)^2 = a_{k+1}^2$. To maximize the Asimov distance, we need to make v_1 small and v_2 large. Apparently, for fixed \mathbf{a} , v_1 is minimized when $b_k = -\text{sign}(a_k)$, which implies $b_i = 0, \forall i \neq k$. To avoid the sign ambiguity, we set $b_k = 1$.

APPENDIX C PROOF OF THEOREM 2

This proof follows the similar steps in the proof of the low-rank case. Denote \mathbb{P} as the subspace spanned by $g_k(\boldsymbol{\Sigma})$ and \mathbb{Q} as the subspace spanned by $g_k(\mathbf{Y})$, and denote their intersection as $\mathbb{T} = \mathbb{P} \cap \mathbb{Q}$. We further denote \mathbf{T} as an orthonormal basis of \mathbb{T} , $[\mathbf{T}, \mathbf{p}]$ as an orthonormal basis of \mathbb{P} , and $[\mathbf{T}, \mathbf{q}]$ as an orthonormal basis of \mathbb{Q} . From the definition of Asimov distance, the subspace distance between \mathbb{P} and \mathbb{Q} is the subspace distance between the span of \mathbf{p} and the span of \mathbf{q} .

First, it is apparent that $\mathbf{q} \in \text{span}[\mathbf{T}, \mathbf{p}, \mathbf{e}_{k+1}, \mathbf{e}_{k+2}, \dots, \mathbf{e}_d]$. Since $\mathbf{q} \perp \mathbf{T}$, we have $\mathbf{q} \in \text{span}[\mathbf{p}, \mathbf{e}]$, where $\mathbf{e} \in \text{span}[\mathbf{e}_{k+1}, \dots, \mathbf{e}_d]$. It is easy to see that the subspace distance between the span of \mathbf{q} and the span of \mathbf{p} will be large if the variance of $\boldsymbol{\Sigma}$ in the span of \mathbf{p} is large and the variance of $\boldsymbol{\Sigma}$ in the span of \mathbf{q} is small. So we should select \mathbf{p} as the direction in $\text{span}[\mathbf{e}_1, \dots, \mathbf{e}_k]$ that has the smallest variance of $\boldsymbol{\Sigma}$ and select \mathbf{e} as the direction among $\text{span}[\mathbf{e}_{k+1}, \dots, \mathbf{e}_n]$ that has the largest variance of $\boldsymbol{\Sigma}$. Since $\mathbf{e} \in \text{span}[\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_d]$ and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k \geq \sigma_{k+1} \geq \dots \geq \sigma_d$, \mathbf{p} should be \mathbf{e}_k and \mathbf{e} should be \mathbf{e}_{k+1} . So, we have $\mathbf{q} \in \text{span}[\mathbf{e}_k, \mathbf{e}_{k+1}]$.

Second, for a fixed direction of \mathbf{a} , let $\hat{\mathbf{a}}$ be the projection of \mathbf{a} onto $\text{span}[\mathbf{e}_k, \mathbf{e}_{k+1}]$. It is easy to see that \mathbf{q} will be closer to \mathbf{a}

as $\hat{\mathbf{a}}$ grows, and as a result, the angle between \mathbf{q} and \mathbf{p} will be larger. This implies the length of \mathbf{a} should be maximized, which indicates $\|\mathbf{a}\| = \eta$ and the distance is maximized when $\mathbf{a} = \hat{\mathbf{a}}$. It also indicates $a_i = 0$ if $i \neq k, k+1$.

Finally, for a fixed \mathbf{a} in the form of (18), the projected variance of \mathbf{Y} in the direction of \mathbf{e}_k is $v_k = \sum_{i \neq k} (a_k b_i)^2 + (a_k b_k + \sigma_k)^2 = a_k^2 + \sigma_k^2 + 2a_k b_k \sigma_k$ and the projected variance of \mathbf{Y} in the direction of \mathbf{e}_{k+1} is $v_{k+1} = \sum_{i \neq k+1} (a_{k+1} b_i)^2 + (\sigma_{k+1} + a_{k+1} b_{k+1})^2 = a_{k+1}^2 + \sigma_{k+1}^2 + 2a_{k+1} b_{k+1}$. To maximize the Asimov distance, we should make v_k small and make v_{k+1} large. With the constraint that $\|\mathbf{b}\| = 1$, we should have $b_k^2 + b_{k+1}^2 = 1$, which implies $b_i = 0$ for all $i \neq k$ and $i \neq (k+1)$.

As shown above, the optimal \mathbf{a} and \mathbf{b} should be in the form of (18) and (19), which completes our proof.

APPENDIX D PROOF OF THEOREM 3

The optimal solution to problem (25) either locates at the boundary, or at the stationary points.

We first characterize the stationary points. At the stationary points, the value (α^*, β^*) satisfies the necessary conditions

$$\begin{cases} \frac{\partial}{\partial \alpha} |\cos \varphi(\alpha, \beta)|_{\alpha=\alpha^*, \beta=\beta^*} = 0, \\ \frac{\partial}{\partial \beta} |\cos \varphi(\alpha, \beta)|_{\alpha=\alpha^*, \beta=\beta^*} = 0. \end{cases} \quad (49)$$

Since $\sin \varphi^* \neq 0$, we have

$$\begin{cases} \frac{\partial}{\partial \alpha} \varphi(\alpha, \beta)|_{\alpha=\alpha^*, \beta=\beta^*} = 0, \\ \frac{\partial}{\partial \beta} \varphi(\alpha, \beta)|_{\alpha=\alpha^*, \beta=\beta^*} = 0, \end{cases} \quad (50)$$

in which

$$\begin{aligned} \frac{\partial \varphi}{\partial \alpha} &= \frac{\eta}{a_x^2 + a_y^2} (\eta(3\sigma_{k+1}^2 - \sigma_k^2 + \eta^2) \\ &\quad + 2\eta(\sigma_k^2 - \sigma_{k+1}^2) \cos^2(\alpha) + 2\eta(\sigma_k^2 - \sigma_{k+1}^2) \cos^2(\beta) \\ &\quad + \sigma_k(\sigma_k^2 - \sigma_{k+1}^2 + 3\eta^2) \cos(\alpha) \cos(\beta) \\ &\quad + \sigma_{k+1}(\sigma_{k+1}^2 - \sigma_k^2 + 3\eta^2) \sin(\alpha) \sin(\beta)), \\ \frac{\partial \varphi}{\partial \beta} &= \frac{\eta}{a_x^2 + a_y^2} (\sigma_k(\sigma_{k+1}^2 + \eta^2 - \sigma_k^2) \sin(\alpha) \sin(\beta) \\ &\quad + \sigma_{k+1}(\sigma_k^2 + \eta^2 - \sigma_{k+1}^2) \cos(\alpha) \cos(\beta) + 2\eta\sigma_k\sigma_{k+1}). \end{aligned}$$

Eliminating $\sin(\alpha) \sin(\beta)$ from (50), we have

$$C \cos^2(\alpha) + D \cos(\alpha) \cos(\beta) + C \cos^2(\beta) + F = 0, \quad (51)$$

where $C = 2\eta\sigma_k(\sigma_k^2 - \sigma_{k+1}^2)(\sigma_{k+1}^2 + \eta^2 - \sigma_k^2)$, $D = (\sigma_k^2 - \sigma_{k+1}^2)(-\sigma_k^2 - \sigma_{k+1}^2 - 2\eta^2(\sigma_k^2 + \sigma_{k+1}^2) + 3\eta^4)$ and $F = \eta\sigma_k(\sigma_k^4 + \sigma_{k+1}^4 + \eta^4 - 2\sigma_k^2\sigma_{k+1}^2 - 2\sigma_k^2\eta^2 - 2\sigma_{k+1}^2\eta^2)$.

Further, we rewrite the first equation of (50) as

$$c\sqrt{(1 - \cos^2(\alpha))(1 - \cos^2(\beta))} + d \cos(\alpha) \cos(\beta) + e = 0, \quad (52)$$

where $c = \sigma_k(\sigma_{k+1}^2 + \eta^2 - \sigma_k^2)$, $d = \sigma_{k+1}(\sigma_k^2 + \eta^2 - \sigma_{k+1}^2)$, and $e = 2\eta\sigma_k\sigma_{k+1}$.

Combining (51) and (52) and eliminating $\cos^2(\alpha)$ and $\cos^2(\beta)$, we have

$$(c^2 - d^2) \cos(\alpha)^2 \cos(\beta)^2 + \left(\frac{Dc^2}{C} - 2de \right) \cos(\alpha) \cos(\beta) + \frac{c^2 F}{C} + c^2 - e^2 = 0.$$

The left side of the equation is a quadratic function with respect to $r = \cos(\alpha) \cos(\beta)$. The two roots are:

$$r_1 = -\frac{\sigma_k \eta}{\sigma_k^2 - \sigma_{k+1}^2}, r_2 = -\frac{\sigma_k}{2} \left(\frac{1}{\eta} + \frac{\eta}{\sigma_k^2 - \sigma_{k+1}^2} \right).$$

Note that $\eta \in [0, \sigma_k - \sigma_{k+1})$, so we have $r_1 \in (-\frac{\sigma_k}{\sigma_k + \sigma_{k+1}}, 0]$, $r_2 \in (-\infty, -\frac{\sigma_k}{\sigma_k - \sigma_{k+1}})$. Since $|\cos(\alpha) \cos(\beta)| \leq 1$, $\frac{\sigma_k}{\sigma_k + \sigma_{k+1}} < 1$, and $\frac{\sigma_k}{\sigma_k - \sigma_{k+1}} > 1$, we should only retain the first root r_1 . Substitute $\cos(\alpha) \cos(\beta) = r_1 = -\frac{\eta \sigma_k}{\sigma_k^2 - \sigma_{k+1}^2}$ into (51), and we have $C \cos^4(\alpha) + (Dr_1 + F) \cos^2(\alpha) + Cr_1^2 = 0$. The left side of the equation is a quadratic function with respect to $s = \cos^2(\alpha)$, so we can easily find its roots. Let us denote s_1 and s_2 as the two roots:

$$s_1 = \frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 - \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)}, s_2 = \frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 + \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)},$$

where $H = \sigma_k^4 + \sigma_{k+1}^4 + \eta^4 - 2\sigma_k^2 \sigma_{k+1}^2 - 2\sigma_k^2 \eta^2 - 2\sigma_{k+1}^2 \eta^2$. We need to check that H is positive. Viewing H as a function of η and taking derivative, we have $H'(\eta) = 2\eta(2\eta^2 - 2(\sigma_k^2 + \sigma_{k+1}^2)) < 0$. Since $\eta^2 \in [0, (\sigma_k - \sigma_{k+1})^2]$, we have $H(\eta) \in (0, (\sigma_k^2 - \sigma_{k+1}^2)^2)$.

As $\cos(\alpha)^2 \leq 1$, we need to check whether $s_1, s_2 \in [0, 1]$.

Firstly, as H is a decreasing function of η in the considered range, s_1 is an increasing function of η . Therefore, we have $\min(s_1) = s_1(\eta)|_{\eta=0} = 0$ and $\max(s_1) = s_1(\eta)|_{\eta=\sigma_k - \sigma_{k+1}} = \frac{\sigma_k}{\sigma_k + \sigma_{k+1}} < 1$. Hence, s_1 is a valid solution.

Secondly, it is easy to check that s_2 is a decreasing function of η . So, we have $\max(s_2) = s_2(\eta)|_{\eta=0} = 1$ and $\min(s_2) = s_2(\eta)|_{\eta=\sigma_k - \sigma_{k+1}} = \frac{\sigma_k}{\sigma_k + \sigma_{k+1}} < 1$, which means s_2 is also a valid solution. Hence, we have two stationary points

$$\begin{cases} \cos^2(\alpha) = \frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 \pm \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)}, \\ \cos^2(\beta) = \frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 \mp \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)}. \end{cases} \quad (53)$$

Since there are two sets of solutions in (53), we should determine which one is better. The variance of \mathbf{Y} in the direction of \mathbf{e}_k is $v_k = \cos^2(\alpha) + \sigma_k^2 + 2 \cos(\alpha) \cos(\beta)$ and the variance of \mathbf{Y} in the direction of \mathbf{e}_{k+1} is $v_{k+1} = \sin^2(\alpha) + \sigma_{k+1}^2 + 2 \sin(\alpha) \sin(\beta)$. Both of the two sets of solutions in (53) lead to $\cos(\alpha) \cos(\beta) = -\frac{\eta \sigma_k}{\sigma_k^2 - \sigma_{k+1}^2}$ and $\sin(\alpha) \sin(\beta) = \frac{\eta \sigma_{k+1}}{\sigma_k^2 - \sigma_{k+1}^2}$. For fixed $\cos(\alpha) \cos(\beta)$ and $\sin(\alpha) \sin(\beta)$, the smaller $\cos^2(\alpha)$ is, the smaller v_k will be, and the larger the subspace distance will be. Hence, we conclude the stationary point that satisfies

$$\begin{cases} \cos^2(\alpha^*) = \frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 - \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)} \\ \cos^2(\beta^*) = \frac{\sigma_k^2 - \sigma_{k+1}^2 + \eta^2 + \sqrt{H}}{2(\sigma_k^2 - \sigma_{k+1}^2)} \end{cases} \quad (54)$$

leads to a larger subspace distance.

Finally, it is easy to compute the objective values of problem (25) at the boundary points. Comparing these values with the objective values induced by the point in equation (54), we can readily conclude the point in equation (54) gives a larger objective value. In summary, given that $\alpha \in [0, \pi/2]$ and $\beta \in [\pi/2, \pi]$, the optimal α and β are shown in (26).

APPENDIX E PROOF OF THEOREM 4

The proof has two main steps. In the first step, we show that non-zero entries of \mathbf{B} are in the k th and $(k+1)$ th rows. In the second step, we will further prove the entries except in the k th and $(k+1)$ th columns should be zero.

In the first step, we follow similar proof procedures in Theorem 2. We use \mathbb{P} to denote the subspace spanned by $g_k(\Sigma)$ and \mathbb{Q} to denote the subspace spanned by $g_k(\mathbf{Y})$. We also use \mathbb{T} to represent the intersection of the two subspaces and further denote \mathbf{T} as one set of orthonormal bases of \mathbb{T} , $[\mathbf{T}, \mathbf{p}]$ as one set of orthonormal bases of \mathbb{P} and $[\mathbf{T}, \mathbf{q}]$ as one set of orthonormal bases of \mathbb{Q} . So, the subspace distance between \mathbb{P} and \mathbb{Q} is the subspace distance between the subspace spanned by \mathbf{p} and that spanned by \mathbf{q} . Following the same arguments in Theorem 2, by setting all the entries of \mathbf{B} to be zero except the k th and $(k+1)$ th rows, we can guarantee achieving the maximal subspace distance and further we have $\mathbf{q} \in \text{span}[\mathbf{e}_k, \mathbf{e}_{k+1}]$ and $\mathbf{p} = \mathbf{e}_k$.

In the second step, since the non-zero elements of \mathbf{B} only locate in the k th and $(k+1)$ th rows and $\mathbf{q} \in \text{span}[\mathbf{e}_k, \mathbf{e}_{k+1}]$, it indicates \mathbf{q} is the direction with the maximal variance on the span of \mathbf{e}_k and \mathbf{e}_{k+1} . Assuming $\mathbf{q} = [0, \dots, \cos(\gamma), \sin(\gamma), \dots, 0]^T$ with $\cos(\gamma)$ and $\sin(\gamma)$ being in the k th and $(k+1)$ th coordinates respectively and according to the definition of principal components, we can find γ by solving the optimization problem

$$\operatorname{argmax}_{\gamma} : \mathbf{q}^T \mathbf{Y} \mathbf{Y}^T \mathbf{q}. \quad (55)$$

Plug $\mathbf{q} = [0, \dots, \cos(\gamma), \sin(\gamma), \dots, 0]^T$ into the objective function, and we have

$$\mathbf{q}^T \mathbf{Y} \mathbf{Y}^T \mathbf{q} = \begin{bmatrix} \cos(\gamma) \\ \sin(\gamma) \end{bmatrix}^T \begin{bmatrix} b_{x1} & \frac{1}{2} b_y \\ \frac{1}{2} b_y & b_{x2} \end{bmatrix} \begin{bmatrix} \cos(\gamma) \\ \sin(\gamma) \end{bmatrix}, \quad (56)$$

where $b_{x1} = \|\mathbf{b}_k + \mathbf{e}_k \sigma_k\|^2$, $b_{x2} = \|\mathbf{b}_{k+1} + \mathbf{e}_{k+1} \sigma_{k+1}\|^2$, $b_y = 2(\mathbf{b}_k + \mathbf{e}_k \sigma_k)^T (\mathbf{b}_{k+1} + \mathbf{e}_{k+1} \sigma_{k+1})$, with \mathbf{b}_k and \mathbf{b}_{k+1} being the transpose of the k th and $(k+1)$ th rows of \mathbf{B} respectively and $\mathbf{e}_k \in \mathbb{R}^n$, $\mathbf{e}_{k+1} \in \mathbb{R}^n$ being the standard bases.

We can solve (56) by computing the first principal component of the middle matrix of the right hand of (56). Using the result from equation (23), we have $\gamma = 0.5 \operatorname{atan}2(b_y, b_x)$, where $b_x = b_{x1} - b_{x2}$. Since the subspace distance is the distance between \mathbf{q} and \mathbf{e}_k , it is apparent that the subspace distance is $|\gamma|$. To maximize $|\gamma|$, we first determine the sign of b_y or b_x . We have

$$\begin{aligned} & \frac{b_x}{\|\mathbf{b}_k + \mathbf{e}_k \sigma_k\| + \|\mathbf{b}_{k+1} + \mathbf{e}_{k+1} \sigma_{k+1}\|} \\ & = \|\mathbf{b}_k + \mathbf{e}_k \sigma_k\| - \|\mathbf{b}_{k+1} + \mathbf{e}_{k+1} \sigma_{k+1}\| \end{aligned}$$

$$\begin{aligned} &\geq \sigma_k - \|\mathbf{b}_k\| - \sigma_{k+1} - \|\mathbf{b}_{k+1}\| \\ &\geq \sigma_k - \sigma_{k+1} - \sqrt{2}\eta \end{aligned} \quad (57)$$

$$> 0, \quad (58)$$

where inequality (57) is the result of the energy constraint that $\eta \geq \|\mathbf{B}\|_F = \sqrt{\|\mathbf{b}_k\|^2 + \|\mathbf{b}_{k+1}\|^2} \geq \frac{1}{\sqrt{2}}(\|\mathbf{b}_k\| + \|\mathbf{b}_{k+1}\|)$, and inequality (58) is due to the assumption that $\eta < \frac{\sigma_k - \sigma_{k+1}}{\sqrt{2}}$. In summary, b_x is positive. Using the property of atan2 function, when $b_x > 0$, maximizing $|\gamma|$ is equivalent to maximizing $|b_y/b_x|$. Thus, we can formulate our problem as

$$\begin{aligned} \max_{\mathbf{b}_k, \mathbf{b}_{k+1}} : & |b_y/b_x| \\ \text{s.t.} & \|[\mathbf{b}_k, \mathbf{b}_{k+1}]\|_F \leq \eta. \end{aligned} \quad (59)$$

In the objective function,

$$\begin{aligned} b_y &= 2(\mathbf{b}_1^\top \mathbf{b}_2 + (b_{k,k} + \sigma_k)b_{k+1,k} \\ &\quad + b_{k,k+1}(b_{k+1,k+1} + \sigma_{k+1})), \\ b_x &= \|\mathbf{b}_1\|^2 - \|\mathbf{b}_2\|^2 + (b_{k,k} + \sigma_k)^2 + b_{k,k+1}^2 - b_{k+1,k}^2 \\ &\quad - (b_{k+1,k+1} + \sigma_{k+1})^2, \end{aligned}$$

where $\mathbf{b}_1 = [b_{k,1}, b_{k,2}, \dots, b_{k,k-1}, b_{k,k+2}, \dots, b_{k,n}]^\top$ and $\mathbf{b}_2 = [b_{k+1,1}, b_{k+1,2}, \dots, b_{k+1,k-1}, b_{k+1,k+2}, \dots, b_{k+1,n}]^\top$ which are the vectors obtained by deleting the k th and $(k+1)$ th elements of \mathbf{b}_k and \mathbf{b}_{k+1} respectively. We can change the sign of b_y/b_x by changing the signs of \mathbf{b}_1 , $b_{k+1,k}$, and $b_{k,k+1}$. Since both of the values b_y/b_x and $-b_y/b_x$ are obtainable, we can remove the absolute value operation. Thus, our objective can be further simplified to maximize b_y/b_x . To complete the proof of Theorem 4, we should further demonstrate that when the optimality of our objective function is obtained, \mathbf{b}_1 and \mathbf{b}_2 should be vectors with all their entries being zero. To prove that, we examine the objective function further

$$\begin{aligned} b_y &\leq 2(\|\mathbf{b}_1\| \|\mathbf{b}_2\| + (b_{k,k} + \sigma_k)b_{k+1,k} \\ &\quad + b_{k,k+1}(b_{k+1,k+1} + \sigma_{k+1})) \\ &\leq 2((b_{k,k} + \sigma_k)b_{k+1,k} \\ &\quad + \sqrt{b_{k,k+1}^2 + \|\mathbf{b}_1\|^2}(\sqrt{b_{k+1,k+1}^2 + \|\mathbf{b}_2\|^2} + \sigma_{k+1})), \end{aligned} \quad (60)$$

$$\begin{aligned} b_x &\geq (b_{k,k} + \sigma_k)^2 + b_{k,k+1}^2 + \|\mathbf{b}_1\|^2 - b_{k+1,k}^2 \\ &\quad - (\sqrt{b_{k+1,k+1}^2 + \|\mathbf{b}_2\|^2} + \sigma_{k+1})^2. \end{aligned} \quad (62)$$

Inequality (60) implies that the optimal value is determined by the norms of \mathbf{b}_1 and \mathbf{b}_2 instead of their specific values. Inequality (61) is true as

$$\begin{aligned} &\sqrt{b_{k,k+1}^2 + \|\mathbf{b}_1\|^2} \left(\sqrt{b_{k+1,k+1}^2 + \|\mathbf{b}_2\|^2} + \sigma_{k+1} \right) \\ &= \sqrt{b_{k,k+1}^2 + \|\mathbf{b}_1\|^2} \sqrt{b_{k+1,k+1}^2 + \|\mathbf{b}_2\|^2} \\ &\quad + \sigma_{k+1} \sqrt{b_{k,k+1}^2 + \|\mathbf{b}_1\|^2} \\ &\geq b_{k,k+1}b_{k+1,k+1} + \|\mathbf{b}_1\| \|\mathbf{b}_2\| + \sigma_{k+1}b_{k,k+1} \\ &= \|\mathbf{b}_1\| \|\mathbf{b}_2\| + b_{k,k+1}(b_{k+1,k+1} + \sigma_{k+1}). \end{aligned}$$

Inequality (62) is due to $-(\sqrt{b_{k+1,k+1}^2 + \|\mathbf{b}_2\|^2} + \sigma_{k+1})^2 \leq -\|\mathbf{b}_2\|^2 - (b_{k+1,k+1} + \sigma_{k+1})^2$. The equalities in (61) and (62) hold when $\|\mathbf{b}_1\| = 0$ and $\|\mathbf{b}_2\| = 0$. This means that, for any feasible solution $(\mathbf{b}_1, \mathbf{b}_2, b_{k,k}, b_{k,k+1}, b_{k+1,k}, b_{k+1,k+1})$ in (59), there is another corresponding feasible solution $(\mathbf{0}, \mathbf{0}, b_{k,k}, \sqrt{b_{k,k+1}^2 + \|\mathbf{b}_1\|^2}, b_{k+1,k}, \sqrt{b_{k+1,k+1}^2 + \|\mathbf{b}_2\|^2})$, which has a larger objective value. In conclusion, \mathbf{b}_1 and \mathbf{b}_2 should be zero vectors when the optimality of (59) is obtained. This completes our proof.

APPENDIX F CONNECTION BETWEEN ASIMOV DISTANCE AND PCR PROBLEM

We first illustrate a connection between the Asimov distance and the projection 2-norm. We then use this connection to establish a connection between the Asimov distance and the PCR problem.

To see the relationship between the Asimov distance and the projection 2-norm, assume \mathbb{X} be the k -dimensional subspace learned from the original data matrix and $\hat{\mathbb{X}}$ be the k -dimensional subspace learned from the modified data matrix. Furthermore, let $\mathbf{P} \in \mathbb{R}^{n \times n}$ be the orthogonal projection onto \mathbb{X} and $\hat{\mathbf{P}} \in \mathbb{R}^{n \times n}$ be the orthogonal projection onto $\hat{\mathbb{X}}$. Then, the Asimov distance between \mathbb{X} and $\hat{\mathbb{X}}$, denoted as $\theta(\mathbb{X}, \hat{\mathbb{X}})$, can also be computed as:

$$\sin \theta = \|\mathbf{P} - \hat{\mathbf{P}}\|_2,$$

where $\|\cdot\|_2$ is the induced 2-norm. Detailed proof can be found in Chapter 2.5 of [18].

Using results in this paper and the aforementioned relationship between Asimov distance and projection 2-norm, we can perform further analysis on the PCR problem. In particular, let $r_1 = \|\mathbf{y} - \mathbf{y}_1\|$ denote the residual after PCR, where \mathbf{y} is the response vector and $\mathbf{y}_1 = \mathbf{P}\mathbf{y}$ is the projection of \mathbf{y} onto the selected k -dimensional subspace according to the original feature matrix. Denote $r_2 = \|\mathbf{y} - \mathbf{y}_2\|$ as the residual of PCR after we modify the feature matrix, where $\mathbf{y}_2 = \hat{\mathbf{P}}\mathbf{y}$ is the projection of \mathbf{y} onto the selected k -dimensional subspace after we modify the feature matrix. The following inequality shows that the difference of the two residuals can be bounded by the product of the norm of \mathbf{y} and the projection 2-norm:

$$\begin{aligned} |r_1 - r_2| &= \|\|\mathbf{y} - \mathbf{y}_1\| - \|\mathbf{y} - \mathbf{y}_2\|\| \\ &= \|\|\mathbf{y} - \mathbf{P}\mathbf{y}\| - \|\mathbf{y} - \hat{\mathbf{P}}\mathbf{y}\|\| \\ &\leq \|(\mathbf{P} - \hat{\mathbf{P}})\mathbf{y}\| \\ &\leq \|\mathbf{P} - \hat{\mathbf{P}}\|_2 \|\mathbf{y}\| \\ &= \|\mathbf{y}\| \sin \theta. \end{aligned}$$

As our analysis shows, θ depends on the energy budget and the singular values of the original feature matrix. Hence, given the energy budget and the original data points, we can establish the largest possible change of the residual compared with the original residual.

REFERENCES

- [1] F. Li, L. Lai, and S. Cui, "On the adversarial robustness of subspace learning," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Brighton, U.K., May 2019, pp. 2477–2481.
- [2] Y. Li, W. Dai, J. Zou, H. Xiong, and Y. F. Zheng, "Structured sparse representation with union of data-driven linear and multilinear subspaces model for compressive video sampling," *IEEE Trans. Signal Process.*, vol. 65, no. 19, pp. 5062–5077, Oct. 2017.
- [3] J. Xin, N. Zheng, and A. Sano, "Subspace-based adaptive method for estimating direction-of-arrival with luenberger observer," *IEEE Trans. Signal Process.*, vol. 59, no. 1, pp. 145–159, Jan. 2011.
- [4] Y. Shen, M. Mardani, and G. B. Giannakis, "Online categorical subspace learning for sketching big data with misses," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4004–4018, Aug. 2017.
- [5] H. Guo, C. Qiu, and N. Vaswani, "An online algorithm for separating sparse and low-dimensional signal sequences from their sum," *IEEE Trans. Signal Process.*, vol. 62, no. 16, pp. 4284–4297, Aug. 2014.
- [6] R. Otazo, E. J. Candès, and D. K. Sodickson, "Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components," *Magn. Reson. Medicine*, vol. 73, no. 3, pp. 1125–1136, Apr. 2015.
- [7] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, Aug. 2009.
- [8] M. Mardani, G. Mateos, and G. B. Giannakis, "Dynamic anomalography: Tracking network anomalies via sparsity and low rank," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 1, pp. 50–66, Feb. 2012.
- [9] H. Guo and N. Vaswani, "Video denoising via online sparse and low-rank matrix decomposition," in *Proc. IEEE Statistical Signal Process. Workshop*, Palma de Mallorca, Spain, Jun. 2016, pp. 1–5.
- [10] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *J. ACM*, vol. 58, no. 3, pp. 11:1–11:37, Jun. 2011.
- [11] D. Hsu, S. M. Kakade, and T. Zhang, "Robust matrix decomposition with sparse corruptions," *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7221–7234, Jun. 2011.
- [12] C. Qiu, N. Vaswani, B. Lois, and L. Hogben, "Recursive robust PCA or recursive sparse recovery in large but structured noise," *IEEE Trans. Inf. Theory*, vol. 60, no. 8, pp. 5007–5039, Jun. 2014.
- [13] Y. Chen, H. Xu, C. Caramanis, and S. Sanghavi, "Robust matrix completion and corrupted columns," in *Proc. Int. Conf. Machine Learn.*, Bellevue, Washington, Jun. 2011, pp. 873–880.
- [14] K. Eykholt *et al.*, "Robust physical-world attacks on deep learning visual classification," in *Proc. IEEE/CVF Conf. Comput. Vision Pattern Recognit.*, Salt Lake City, UT, Jun. 2018, pp. 1625–1634.
- [15] N. Carlini *et al.*, "Hidden voice commands," in *Proc. USENIX Secur. Symp.*, Austin, TX, Aug. 2016, pp. 513–530.
- [16] S. G. Finlayson, J. D. Bowers, J. Ito, J. L. Zittrain, A. L. Beam, and I. S. Kohane, "Adversarial attacks on medical machine learning," *Science*, vol. 363, no. 6433, pp. 1287–1289, Mar. 2019.
- [17] X. Chen, C. Liu, B. Li, K. Lu, and D. Song, "Targeted backdoor attacks on deep learning systems using data poisoning," Dec. 2017, *arXiv:1712.05526*.
- [18] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD, USA: The Johns Hopkins Univ. Press, 2013.
- [19] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM J. Matrix Anal. Appl.*, vol. 20, no. 2, pp. 303–353, Apr. 1998.
- [20] A. Weinstein, "Almost invariant submanifolds for compact group actions," *J. Eur. Math. Soc.*, vol. 2, no. 1, pp. 53–86, Mar. 2000.
- [21] T. T. Georgiou and M. C. Smith, "Optimal robustness in the gap metric," *IEEE Trans. Autom. Control*, vol. 35, no. 6, pp. 673–686, Jun. 1990.
- [22] G. Vinnicombe, "Frequency domain uncertainty and the graph topology," *IEEE Trans. Autom. Control*, vol. 38, no. 9, pp. 1371–1383, Sep. 1993.
- [23] L. Qui and E. Davison, "Feedback stability under simultaneous gap metric uncertainties in plant and controller," *Syst. Control Lett.*, vol. 18, no. 1, pp. 9–22, Jan. 1992.
- [24] C. He and J. M. Moura, "Robust detection with the gap metric," *IEEE Trans. Signal Process.*, vol. 45, no. 6, pp. 1591–1604, Jun. 1997.
- [25] P. A. Absil, A. Edelman, and P. Koev, "On the largest principal angle between random subspaces," *Linear Algebra Appl.*, vol. 414, no. 1, pp. 288–294, Apr. 2006.
- [26] M. Jagielski, A. Oprea, B. Biggio, C. Liu, C. Nita-Rotaru, and B. Li, "Manipulating machine learning: Poisoning attacks and countermeasures for regression learning," in *Proc. IEEE Symp. Secur. Privacy*, San Francisco, CA, May 2018, pp. 19–35.
- [27] E. Bayraktar and L. Lai, "On the adversarial robustness of multivariate robust estimation," 2019, *arXiv:1903.11220*.
- [28] D. L. Pimentel-Alarcón, A. Biswas, and C. R. Solís-Lemus, "Adversarial principal component analysis," in *Proc. IEEE Int. Symp. Inf. Theory*, Aachen, Germany, Jun. 2017, pp. 2363–2367.
- [29] R. Zimmermann, "A geometric note on subspace updates and orthogonal matrix decompositions under rank-one modifications," Mar. 2019, *arXiv:1711.08235*.
- [30] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 2012.
- [31] R. C. Thompson, "The behavior of eigenvalues and singular values under perturbations of restricted rank," *Linear Algebra Appl.*, vol. 13, no. 1-2, pp. 69–78, 1976.
- [32] J. E. Jackson, *A User's Guide to Principal Components*, vol. 587. Hoboken, NJ, USA: Wiley, 2005.
- [33] J. H. Kalivas, "Two data sets of near infrared spectra," *Chemometrics Intell. Laboratory Syst.*, vol. 37, no. 2, pp. 255–259, Jun. 1997.