LBNL-43597
Preprint

ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

Structure and Perturbation Analysis of Truncated SVD for Column-Partitioned Matrices

Zhenyue Zhang and Hongyuan Zha

National Energy Research Scientific Computing Division

August 1999

Submitted to
Siam Journal of Matrix Analysis

REFERENCE COPY
Does Not
Circulate
Bldg. 50 Library - Ref.
Lawrence Berkeley National Laboratory

LBNL-43597
Copy 1

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBNL-43597

Structure and Perturbation Analysis of Truncated SVD for Column-Partitioned Matrices

Zhenyue Zhang and Hongyuan Zha

National Energy Research Scientific Computing Division
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, California 94720

August 1999

This work was supported in part by the Director, Office of Laboratory Policy and Infrastructure Management, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098, the Natural Science Foundation of China under Project No. 19771073, and the National Science Foundation under Grant No. CCR-9619452.

STRUCTURE AND PERTURBATION ANALYSIS OF TRUNCATED SVD FOR COLUMN-PARTITIONED MATRICES

ZHENYUE ZHANG* AND HONGYUAN ZHA†

Abstract. We present a detailed study of truncated SVD for column-partitioned matrices. In particular, we analyze the relation between the truncated SVD of a matrix and the truncated SVDs of its submatrices. We give necessary and sufficient conditions under which truncated SVD of a matrix can be constructed from those of its submatrices. We also present perturbation analysis to show that an *approximate* truncated SVD can still be computed even if the given necessary and sufficient conditions are only approximately satisfied.

1. Introduction. In many applications, it is desirable to compute a low-rank approximation of a given matrix $A \in \mathcal{R}^{m\times n}$, and the matrix A can be large and/or sparse, see [5], for example, for a list of application areas. The theory of singular value decomposition (SVD) provides the following characterization of the best low-rank approximation of A in terms of Frobenius norm $\|\cdot\|_F$ [3]. (Similar results hold for general unitarily-invariant norms.)

THEOREM 1.1. *Let the SVD of $A \in \mathcal{R}^{m\times n}$ be $A = U\Sigma V^T$ with*

$$\Sigma = \operatorname{diag}(\sigma_1, \ldots, \sigma_{\min(m,n)}), \quad \sigma_1 \geq \cdots \geq \sigma_{\min(m,n)},$$

and U and V orthogonal. Then for $1 \leq k \leq \min(m,n)$,

$$\sum_{i=k+1}^{\min(m,n)} \sigma_i^2 = \min\{\, \|A - B\|_F^2 \mid \operatorname{rank}(B) \leq k\}.$$

And the minimum is achieved with $\operatorname{best}_k(A) \equiv U_k \operatorname{diag}(\sigma_1, \ldots, \sigma_k) V_k^T$, *where U_k and V_k are the matrices formed by the first k columns of U and V, respectively. Furthermore,* $\operatorname{best}_k(A)$ *is unique if and only if $\sigma_k > \sigma_{k+1}$.*

In this paper, we call $\operatorname{best}_k(A)$ a *truncated SVD* of A, which is obtained by truncating the expansion

$$A = \sum_{i=1}^{\min(m,n)} \sigma_i u_i v_i^T$$

up to the kth term. Here $U = [u_1, \ldots, u_m]$ and $V = [v_1, \ldots, v_n]$. Algorithms for computing (truncated) SVD, even in the case when A is large and sparse are well established [1, 2, 3]. In this paper we are concerned with an interesting issue which is motivated by some of the results developed in [12] where we dealt with the relation of truncated SVD and a special indexing method *latent semantic indexing* used

* Center for Mathematical Sciences & Department of Applied Mathematics, Zhejiang University, Hangzhou, 310027, P. R. China. zyzhang@math.zju.edu.cn. The work of this author was supported in part by NSFC (project 19771073), Zhejiang Provincial Natural Science Foundation of China, Scientific Research Foundation for Returned Overseas Chinese Scholars, State Education Commission, NSF grant CCR-9619452, the Director, Office of Science, Office of Laboratory Policy and Infrastructure Management, of the U.S. Department of Energy under Contract No. DE-AC-76SF00098. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy.

† Department of Computer Science and Engineering, The Pennsylvania State University, University Park, PA 16802, zha@cse.psu.edu. The work of this author was supported in part by NSF grant CCR-9619452.

in information retrieval. We will build on the results obtained in [12] and study truncated SVD of column-partitioned matrices in greater generality. We observed that in some applications, the A is naturally partitioned into several block columns: $A = [A_1, \ldots, A_s]$. In text categorization applications, for example, each column of A represents a document in a given text corpus, and A_i consists of all the documents in the text corpus that are about a particular topic i. In dynamic information retrieval applications, A_1 can be the documents from an old text corpus, and $A_2, \ldots, A_s$ are document collections added dynamically as new documents become available [9]. An important problem from those applications is the following: we have computed the truncated SVD of some of the A_i's, say, $\text{best}_k[A_1, \ldots, A_t], t < s$, and the matrix $[A_1, \ldots, A_t]$ has been discarded and is therefore no longer available. How can we construct a truncated SVD of A from $\text{best}_k[A_1, \ldots, A_t]$ and the remaining $[A_{t+1}, \ldots, A_s]$? To answer this question we need to study the relation between the truncated SVD of a matrix and the truncated SVDs of its submatrices. It turns out that a general theory can be developed and the questions we are interested in can be answered by certain special cases of the general theory.

The rest of the paper is organized as follows: In Section 2, we give necessary and sufficient conditions that guarantee a truncated SVD of a column-partitioned matrix A can be perfectly constructed from truncated SVD's of its submatrices. The orthogonality of certain submatrices of A plays an important role in specifying those conditions. We also relate the sufficient conditions to a class of matrices with the so-called low-rank-plus-shift structure [6, 7, 10]. In Section 3, we expand the results in Section 2 to the case where the necessary and sufficient conditions are only approximately satisfied by the given matrix A. We show that a truncated SVD of A can be approximately constructed from truncated SVD's of its submatrices. Along the way, we prove some novel perturbation bounds for truncated SVD of a matrix that are of their own interests. The case for matrices with low-rank-plus-shift structure is analyzed in some detail, and an improved perturbation bound is also derived.

2. Necessary and Sufficient Conditions. As mentioned in Section 1, we are interested in finding conditions on a column-partitioned matrix $A = [A_1, \ldots, A_s]$ such that a truncated SVD of A can be constructed from those of the A_i's. At first glance, using only truncated SVD's of the A_i's certainly loses some information about the original matrix A. Therefore, in general, we can not expect to reconstruct a truncated SVD of A *perfectly* from those of the A_i's. The goal of this section is to find conditions under which this can be done. We first present a general result which gives the necessary and sufficient condition for a matrix and its perturbation to have the same truncated SVD's.

NOTE. Throughout the rest of the paper, we will use the following convention: whenever $\text{best}_k(B)$ is mentioned for a matrix B, it is implicitly assumed that $\sigma_k(B) > \sigma_{k+1}(B)$ so that $\text{best}_k(B)$ is uniquely defined.

THEOREM 2.1. *Let $A = B + C$. Then $\text{best}_k(A) = \text{best}_k(B)$ if and only if*

$$C^T \text{best}_k(B) = 0, \quad \text{best}_k(B) C^T = 0, \quad \sigma_k(B) > \sigma_{k+1}(A).$$

Proof. We first deal with the "only if" part of the proof which is rather straightforward. Since

$$(A - \text{best}_k(A))^T \text{best}_k(A) = 0,$$

it follows from $\text{best}_k(A) = \text{best}_k(B)$ that

$$(A - \text{best}_k(B))^T \text{best}_k(B) = 0.$$

Substituting A with $B + C$ and using the equality $(B - \text{best}_k(B))^T \text{best}_k(B) = 0$, we obtain

$$C^T \text{best}_k(B) = 0.$$

We can similar show $\text{best}_k(B)C^T = 0$. Also $\sigma_k(B) > \sigma_{k+1}(A)$ follows from $\sigma_k(A) = \sigma_k(B)$ and $\sigma_k(A) > \sigma_{k+1}(A)$.

Now we prove the "if" part. Let the SVD of B and C be

$$B = [U_1, U_2] \begin{bmatrix} \Sigma_1 & \\ & \Sigma_2 \end{bmatrix} [V_1, V_2]^T, \quad C = QDG^T,$$

respectively, where $\Sigma_1 \in \mathcal{R}^{k\times k}$ and the matrices are partitioned conformally. Then $\text{best}_k(B) = U_1\Sigma_1 V_1$. Now the conditions $C^T \text{best}_k(B) = 0$ and $\text{best}_k(B)C^T = 0$ implies

$$U_1^T C = 0, \quad CV_1 = 0.$$

Let the SVD of $U_2\Sigma_2 V_2^T + C$ be

$$U_2\Sigma_2 V_2^T + C = \tilde{U}_2\tilde{\Sigma}_2\tilde{V}_2^T.$$

It is readily verified that $\tilde{U}_2^T U_1 = 0$, and $\tilde{V}_2^T V_1 = 0$. Therefore,

$$A = B + C = [U_1, \tilde{U}_2] \begin{bmatrix} \Sigma_1 & \\ & \tilde{\Sigma}_2 \end{bmatrix} [V_1, \tilde{V}_2]^T$$

gives the SVD of A. Since $\sigma_{\min}(\Sigma_1) = \sigma_k(B) > \sigma_{k+1}(A)$, it follows that $\sigma_{\min}(\Sigma_1) > \sigma_{\max}(\tilde{\Sigma}_2)$, and therefore

$$\text{best}_k(A) = U_1\Sigma_1 V_1 = \text{best}_k(B),$$

completing the proof. □

REMARK. We notice that the condition $\sigma_k(B) > \sigma_{k+1}(A)$ does not impose an upper bound on the norm of the perturbation matrix C. Even when $k = 1$, for certain C with norm as large as possible the condition can still be satisfied.

With the above general result, let us now consider A partitioned in various forms. First we partition A as $A = [A_1, A_2]$, where $A_i \in \mathcal{R}^{m\times n_i}, i = 1, 2$. To apply the result of Theorem 2.1, we will write A as the sum of two matrices. For example,

$$A = [A_1, 0] + [0, A_2], \quad A = [\text{best}_{k_1}(A_1), 0] + [A_1 - \text{best}_{k_1}(A_1), A_2],$$

and so on. With these kinds of partitions, the proof of the following corollaries follows straightforwardly from Theorem 2.1, and therefore is omitted here.

COROLLARY 2.2. *Let* $A = [A_1, A_2]$. *Then*

$$\text{best}_k(A) = \text{best}_k([A_1, 0])$$

if and only if

$$A_2^T \text{best}_k(A_1) = 0, \quad \sigma_k(A_1) > \sigma_k(A).$$

COROLLARY 2.3. *Let* $A = [A_1, A_2]$ *and* $k_1 \le n_1$. *Then*

$$\mathrm{best}_k(A) = \mathrm{best}_k([\mathrm{best}_{k_1}(A_1), A_2])$$

if and only if

$$(A_1 - \mathrm{best}_{k_1}(A_1))^T \mathrm{best}_k([\mathrm{best}_{k_1}(A_1), A_2]) = 0, \quad \sigma_k([\mathrm{best}_{k_1}(A_1), A_2]) > \sigma_{k+1}(A).$$

COROLLARY 2.4. *Let* $A = [A_1, A_2]$ *and* $k_1 \le n_1$ *and* $k_2 \le n_2$. *Then*

$$\mathrm{best}_k(A) = \mathrm{best}_k([\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_2)])$$

if and only if

$$[A_1 - \mathrm{best}_{k_1}(A_1), A_2 - \mathrm{best}_{k_2}(A_2)]^T \mathrm{best}_k([\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_2)]) = 0,$$
$$\sigma_k([\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_2)]) > \sigma_{k+1}(A).$$

REMARK. The conditions listed in Corollary 2.4 seem to be rather complicated, however, in some situations, we may be able to verify some stronger but simpler conditions. For example, the following two equalities

$$(A_1 - \mathrm{best}_{k_1}(A_1))^T A_2 = 0, \quad (A_2 - \mathrm{best}_{k_2}(A_2))^T A_1 = 0$$

imply the condition

$$[A_1 - \mathrm{best}_{k_1}(A_1), A_2 - \mathrm{best}_{k_1}(A_2)]^T \mathrm{best}_k([\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_2)]) = 0.$$

An example of this is given in Theorem 2.7.

Now we show another interesting application of Corollary 2.4.

COROLLARY 2.5. *The equality* $\mathrm{best}_k(A) = \mathrm{best}_k([\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_2)])$ *holds if and only if for any* $t_i \ge k_i, i = 1, 2,$

$$\mathrm{best}_k(A) = \mathrm{best}_k([\mathrm{best}_{t_1}(A_1), \mathrm{best}_{t_2}(A_2)]).$$

Proof. We just need to prove the "only if" part. Let $\tilde{A}_i = \mathrm{best}_{t_i}(A_i), i = 1, 2$. It is easy to verify that $\mathrm{best}_{k_i}(\tilde{A}_i) = \mathrm{best}_{k_i}(A_i)$. Now we only need to prove that

$$\mathrm{best}_k([\tilde{A}_1, \tilde{A}_2]) = \mathrm{best}_k([\mathrm{best}_{k_1}(\tilde{A}_1), \mathrm{best}_{k_2}(\tilde{A}_2)]).$$

Using Corollary 2.4, we need to first verify that

$$[\tilde{A}_1 - \mathrm{best}_{k_1}(\tilde{A}_1), \tilde{A}_2 - \mathrm{best}_{k_1}(\tilde{A}_2)]^T \mathrm{best}_k([\mathrm{best}_{k_1}(\tilde{A}_1), \mathrm{best}_{k_2}(\tilde{A}_2)]) = 0.$$

Since $\mathrm{span}\{\tilde{A}_i - \mathrm{best}_{k_i}(\tilde{A}_i)\} \subset \mathrm{span}\{A_i - \mathrm{best}_{k_i}(A_i)\}$, the above equality follows from the given condition. Next the inequality

$$\sigma_k([\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_2)]) \le \sigma_k([\mathrm{best}_{t_1}(A_1), \mathrm{best}_{t_2}(A_2)])$$

follows from a general inequality about the monotonicity of singular values established in [8]. □

The results in Corollaries 2.4 and 2.5 can be generalized to the cases where $A = [A_1, \ldots, A_s]$. We just state the case for Corollary 2.4.

COROLLARY 2.6. *Let* $A = [A_1, \ldots, A_s]$ *with* $A_i \in \mathcal{R}^{m\times n_i}$, *and* $k_i \leq n_i, i = 1, \ldots, s$. *Then*

$$\text{best}_k(A) = \text{best}_k([\text{best}_{k_1}(A_1), \ldots, \text{best}_{k_s}(A_s)])$$

if and only if for $i = 1, \ldots, s$

$$(A_i - \text{best}_{k_i}(A_i))^T \text{best}_k([\text{best}_{k_1}(A_1), \ldots, \text{best}_{k_s}(A_s)]) = 0,$$

and

$$\sigma_k([\text{best}_{k_1}(A_1), \ldots, \text{best}_{k_s}(A_s)]) > \sigma_{k+1}(A).$$

Matrices with Low-Rank-Plus-Shift Structure. As an application of the results established in the above corollaries, we consider a special class of matrices that possess the so-called low-rank-plus-shift structure. This kind of matrices arises naturally in applications such as array signal processing and Latent Semantic Indexing in information retrieval [6, 7, 10]. Specifically, a matrix has the low-rank-plus-shift structure if its cross-product is a low-rank perturbation of a positive multiple of the identity matrix (cf. Equation (2.1)). We now show that matrices with low-rank-plus-shift structure satisfies the sufficient conditions of Corollary 2.4.[1]

THEOREM 2.7. *Let* $A = [A_1, A_2] \in \mathcal{R}^{m\times n}$ *with* $A_1 \in \mathcal{R}^{m\times n_1}$ *and* $A_2 \in \mathcal{R}^{m\times n_2}$. *Assume that*

$$A^T A = X + \sigma^2 I, \tag{2.1}$$

where X *is positive semi-definite with* $\text{rank}(A) = k$. *Partition* X *as* $X = (X_{ij})_{i,j=1}^2$ *with* $X_{ii} \in \mathcal{R}^{n_i\times n_i}$ *and let* $\text{rank}(X_{ii}) = k_i$, $i = 1, 2$. *Then*

$$(A_1 - \text{best}_{k_1}(A_1))^T A_2 = 0, \quad (A_2 - \text{best}_{k_2}(A_2))^T A_1 = 0, \tag{2.2}$$

and furthermore,

$$\text{best}_k(A) = \text{best}_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]).$$

Proof. For $i = 1, 2$, we have $A_i^T A_i = X_{ii} + \sigma^2 I$ with X_{ii} positive semi-definite, and $\text{rank}(X_{ii}) = k_i$. We can write the SVD of A_i in the following form

$$A_i = U_i \,\text{diag}(\Sigma_i, \sigma^2 I) V_i^T = [U_{i1}, U_{i2}]\,\text{diag}(\Sigma_i, \sigma^2 I)[V_{i1}, V_{i2}]^T,$$

where V_i is orthogonal, and

$$\Sigma_i = (D_i + \sigma^2 I)^{1/2}, \quad D_i = \text{diag}(\mu_1^{(i)}, \ldots, \mu_{k_1}^{(i)})$$

with $\mu_1^{(i)} \geq \ldots \geq \mu_{k_1}^{(i)} > 0$. Hence

$$\text{best}_{k_i}(A_i) = U_{i1}\Sigma_i V_{i1}^T, \quad A_i - \text{best}_{k_i}(A_i) = \sigma U_{i2} V_{i2}^T,$$

[1] A similar result was also proved in [12].

and we only need to show $U_{12}^T A_2 = 0$, and $U_{22}^T A_1 = 0$. To this end, consider the symmetric positive semi-definite matrix

$$\begin{bmatrix} V_1 & \\ & V_2 \end{bmatrix}^T (A^T A - \sigma^2 I) \begin{bmatrix} V_1 & \\ & V_2 \end{bmatrix} = \begin{bmatrix} D_1 & 0 & \Sigma_1 U_{11}^T U_{21} \Sigma_2 & \Sigma_1 U_{11}^T U_{22} \Sigma_2 \\ & 0 & \Sigma_1 U_{12}^T U_{21} \Sigma_2 & \Sigma_1 U_{12}^T U_{22} \Sigma_2 \\ & & D_2 & 0 \\ & & & 0 \end{bmatrix},$$

where for the last matrix in the above equation, *blank* denotes element by symmetry. Since a principal submatrix of positive semi-definite is still positive semi-definite, we obtain

$$U_{12}^T U_{21} = 0, \quad U_{11}^T U_{22} = 0, \quad U_{12}^T U_{22} = 0,$$

and the rank of the matrix

$$\tilde{A} = \begin{bmatrix} D_1 & \Sigma_1 U_{11}^T U_{21} \Sigma_2 \\ (\Sigma_1 U_{11}^T U_{21} \Sigma_2)^T & D_2 \end{bmatrix}$$

equals k. Hence by

$$B^T B = \operatorname{diag}(V_{11}, V_{21})(\tilde{A} + \sigma^2 I) \operatorname{diag}(V_{11}^T, V_{21}^T),$$

where $B = [\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]$, we obtain

$$\sigma_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]) > \sigma = \sigma_{k+1}(A).$$

The result of the theorem now follows from Corollary 2.4. □

REMARK. The above results can also be generalized to the cases where $A = [A_1, \ldots, A_s]$.

REMARK. By definition $k_1 \le k$, $k_2 \le k$ and $k \le k_1 + k_2 \le 2k$, and it is easy to find examples for which $k_1 + k_2 = k$ or $k_1 + k_2 = 2k$. In some cases, it is possible to find a permutation P such that $AP = [A_1, A_2]$ will have A_i with k_i, $i = 1, 2$, that are smaller than those of A. For example,

$$A = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}.$$

It is easy to verify that $k = 2$. If we take the first two columns as A_1 and the last two columns as A_2, we have $k_1 = k_2 = 2$. However, if we take the middle two columns as A_1 and the first and last tow columns as A_2, we have $k_1 = k_2 = 1$. This example motivates the following questions: is it possible to find a permutation P such that a partition of $AP = [A_1, A_2]$ with A_1 and A_2 having about the same column dimensions will give $k_1 + k_2 < 2k$? The answer turns out to be no. In the following we show that we can find a class of matrices A satisfying

$$A^T A - \sigma^2 I = X$$

with X positive semi-definite such that for any permutation $AP = [A_1, A_2]$ we will have $k_1 = k_2 = k$, provided the column dimensions of A_i and A_2 are no smaller than k. Let $Y \in \mathcal{R}^{k \times n}$ be a matrix any k columns of which are linearly independent. Let

$$C^T C = Y^T Y + \sigma^2 I$$

be the Cholesky decomposition of $Y^T Y + \sigma^2 I$. Set $A = QC$, where Q is arbitrary orthogonal matrix. Then it is easy to see that for any permutation P, a partition of $AP = [A_1, A_2]$ with column dimensions of A_1 and A_2 at least k will have $k_1 = k_2 = k$.

3. Perturbation Analysis. In the previous section we give necessary and sufficient conditions to perfectly reconstruct a truncated SVD of a matrix from those of its submatrices. In this section, we consider the case when these conditions are no longer satisfied. We first give a general result concerning perturbation bounds of truncated SVD. The perturbation bound is so derived such that we get back the result of Theorem 2.1 when the necessary and sufficient conditions of Theorem 2.1 are satisfied.

THEOREM 3.1. *Let $A = B + C$ with $\sigma_{k+1}(A) < \sigma_k(B)$. Then*

$$\|\mathrm{best}_k(A) - \mathrm{best}_k(B)\| \le \frac{\|A\|_2(\|\mathrm{best}_k(B)C^T\| + \|C^T\mathrm{best}_k(B)\|)}{\sigma_k^2(B) - \sigma_{k+1}^2(A)} + \|P_{\mathrm{best}_k(B)}C^T\|,$$

where $P_{\mathrm{best}_k(B)}$ is the orthogonal projector onto the subspace $\mathrm{span}\{\mathrm{best}_k(B)\}$.

Proof. Let the SVD of B be

$$B = U\Sigma V^T = [U_1, U_2]\,\mathrm{diag}(\Sigma_1, \Sigma_2)[V_1, V_2]^T$$

with $\Sigma_1 \in \mathcal{R}^{k\times k}$ and $\mathrm{best}_k(B) = U_1\Sigma_1 V_1^T$. Write

$$\tilde{C} \equiv U^T C V = \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix}$$

with $C_{11} \in \mathcal{R}^{k\times k}$, and let

$$U^T A V = \Sigma + \tilde{C} = QDG^T = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix} \begin{bmatrix} D_1 & \\ & D_2 \end{bmatrix} \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix}^T \tag{3.3}$$

be the SVD of $U^T A V$ with Q_{11}, D_1 and G_{11} all k-by-k matrices. Then we have

$$\begin{aligned} \|\Delta\| &= \|\mathrm{best}_k(A) - \mathrm{best}_k(B)\| \\ &= \left\| \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix} D_1 \begin{bmatrix} G_{11} \\ G_{21} \end{bmatrix}^T - \begin{bmatrix} \Sigma_1 & \\ & 0 \end{bmatrix} \right\| \\ &= \left\| \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix} [D_1, 0] - \begin{bmatrix} \Sigma_1 \\ 0 \end{bmatrix} [G_{11}, G_{12}] \right\| \\ &= \left\| \begin{bmatrix} Q_{11}D_1 - \Sigma_1 G_{11} & -\Sigma_1 G_{12} \\ Q_{21}D_1 & 0 \end{bmatrix} \right\|. \end{aligned}$$

From (3.3) we have

$$\Sigma_1 G_{11} + [C_{11}, C_{12}] \begin{bmatrix} G_{11} \\ G_{21} \end{bmatrix} = Q_{11}D_1,$$

$$\Sigma_2 G_{21} + [C_{21}, C_{22}] \begin{bmatrix} G_{11} \\ G_{21} \end{bmatrix} = Q_{21}D_1.$$

It follows that

$$\begin{aligned} \|\Delta\| &= \left\| \begin{bmatrix} C_{11}G_{11} + C_{12}G_{21} & -\Sigma_1 G_{12} \\ C_{21}G_{11} + (\Sigma_2 + C_{22})G_{21} & 0 \end{bmatrix} \right\| \\ &= \left\| A \begin{bmatrix} 0 & -G_{12} \\ G_{21} & 0 \end{bmatrix} + \begin{bmatrix} C_{11} \\ C_{21} \end{bmatrix} [G_{11}, G_{12}] \right\| \qquad (3.4) \\ &\le \|A\|\|G_{12}\| + \left\| \begin{bmatrix} C_{11} \\ C_{21} \end{bmatrix} \right\|, \end{aligned}$$

where we have used $\|G_{12}\| = \|G_{21}\|$. On the other hand from the equations

$$(\Sigma + \tilde{C})G = QD, \quad (\Sigma + \tilde{C})Q = GD,$$

we obtain

$$\Sigma_1 G_{12} + [C_{11}, C_{12}] \begin{bmatrix} G_{12} \\ G_{22} \end{bmatrix} = Q_{12} D_2,$$

$$\Sigma_1 Q_{12} + [C_{11}^T, C_{12}^T] \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix} = G_{12} D_2.$$

Therefore,

$$\Sigma_1^2 G_{12} + \Sigma_1 [C_{11}, C_{12}] \begin{bmatrix} G_{11} \\ G_{21} \end{bmatrix} = \Sigma_1 Q_{12} D_2 = G_{12} D_2^2 - [C_{11}^T, C_{12}^T] \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix} D_2,$$

and we have

$$\begin{aligned} (\sigma_k^2(B) - \sigma_{k+1}^2(A))\|G_{12}\| &\le \|\Sigma_1^2 G_{12} - G_{12} D_2^2\| \\ &\le \|[C_{11}^T, C_{12}^T]\|\|D_2\| + \|\Sigma_1 [C_{11}, C_{12}].\| \end{aligned}$$

Recall that $\|D_2\| = \sigma_{k+1}(A) < \sigma_k(B)$. Furthermore,

$$\|\Sigma_1 [C_{11}, C_{12}]\| = \|\Sigma_1 U_1^T C V\| = \|V_1^T (\mathrm{best}_k(B))^T C\| \le \|C^T \mathrm{best}_k(B)\|$$

and

$$\|[C_{11}^T, C_{12}^T]\| = \|V_1^T C^T U^T\| = \|V_1^T C^T\| = \|P_{\mathrm{best}_k(B^T)} C^T\| \le \|\mathrm{best}_k(B) C^T\| / \sigma_k(B).$$

We obtain

$$\|G_{12}\| \le (\|\mathrm{best}_k(B) C^T\| + \|C^T \mathrm{best}_k(B)\|)/(\sigma_k^2(B) - \sigma_{k+1}^2(A)).$$

Substituting the above into Equation (3.4) completing the proof. □

REMARK. As mentioned before the condition $\sigma_{k+1}(A) < \sigma_k(B)$ does not imply an upper bound on the norm of C, i.e., certain C with large norm can still produce a small perturbation of $\mathrm{best}_k(B)$.

If we ignore the structural relationship between B and C, we can derive the following less sharp result.

COROLLARY 3.2. *Under the same condition of Theorem 3.1, we have*

$$\|\mathrm{best}_k(A) - \mathrm{best}_k(B)\| \le \left(1 + \frac{(\|B\| + \|A\|)\|B\|}{\sigma_k^2(B) - \sigma_{k+1}^2(A)}\right) \|C\|.$$

REMARK. If we write the SVD of A as

$$A = \tilde{U} \tilde{\Sigma} \tilde{V}^T = [\tilde{U}_1, \tilde{U}_2] \,\mathrm{diag}(\tilde{\Sigma}_1, \tilde{\Sigma}_2) [\tilde{V}_1, \tilde{V}_2]^T.$$

It is easy to verify that

$$\begin{aligned} \mathrm{best}_k(B) &= U_1 \Sigma_1 V_1^T = B P_{V_1}, \\ \mathrm{best}_k(A) &= \tilde{U}_1 \tilde{\Sigma}_1 \tilde{V}_1^T = A P_{\tilde{V}_1}, \end{aligned}$$

where $P_{V_1} = V_1 V_1^T$ and $P_{\tilde{V}_1} = \tilde{V}_1 \tilde{V}_1^T$ are the orthogonal projectors onto the subspaces spanV_1 and span$\tilde{V}_1$, respectively. Then we can use $\sin\Theta$ Theorem to obtain the bound in Corollary 3.2 [4, Theorem 11.7.2].

COROLLARY 3.3. *Under the condition of Theorem 3.1, if furthermore, we have* $BC^T = 0$, *then*

$$\|\text{best}_k(A) - \text{best}_k(B)\| \le \frac{\|A\|}{\sigma_k^2(B) - \sigma_{k+1}^2(A)} \|C^T \text{best}_k(B)\|.$$

Proof. It is easy to see that $BC^T = 0$ implies $\text{best}_k(B)C^T = 0$ and $P_{\text{best}_k(B)}C^T = 0$. Then the result directly follows from Theorem 3.1. □

Now as we did in Section 2, we now consider A arranged in various forms. We partition A as $A = [A_1, A_2]$, and split A as the sum of two matrices $A = B + C$, for example,

$$A = [A_1, 0] + [0, A_2], \quad A = [\text{best}_{k_1}(A_1), 0] + [A_1 - \text{best}_{k_1}(A_1), A_2],$$

and so on. It is easily verified that for all the splittings used in the following three corollaries, we always have $BC^T = 0$. Then the following perturbation results follow directly from Corollary 3.3.

COROLLARY 3.4. *Let* $A = [A_1, A_2]$ *and* $\sigma_k(A_1) > \sigma_{k+1}(A)$. *Then*

$$\|\text{best}_k(A) - \text{best}_k([A_1, 0])\| \le \frac{\|A\|}{\sigma_k^2(A_1) - \sigma_{k+1}^2(A)} \|A_2^T \text{best}_k(A_1)\|$$

COROLLARY 3.5. *Let* $A = [A_1, A_2]$ *with* $\sigma_k([\text{best}_{k_1}(A_1), A_2]) > \sigma_{k+1}(A)$. *Then*

$$\|\text{best}_k(A) - \text{best}_k([\text{best}_{k_1}(A_1), A_2])\| \le \frac{\|A\|\eta}{\sigma_k^2([\text{best}_{k_1}(A_1), A_2]) - \sigma_{k+1}^2(A)},$$

where

$$\eta = \|(A_1 - \text{best}_{k_1}(A_1))^T \text{best}_k([\text{best}_{k_1}(A_1), A_2])\| \le \|(A_1 - \text{best}_{k_1}(A_1))^T A_2\|.$$

COROLLARY 3.6. *Let* $A = [A_1, A_2]$. *If* $\sigma_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]) > \sigma_k(A)$, *then*

$$\|\text{best}_k(A) - \text{best}_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)])\| \le \frac{\|A\|\eta}{\sigma_k^2([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]) - \sigma_{k+1}^2(A)},$$

where

$$\begin{aligned}\eta &= \|[A_1 - \text{best}_{k_1}(A_1), A_2 - \text{best}_{k_2}(A_2)]^T \text{best}_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)])\| \\ &\le \max\{\|(A_1 - \text{best}_{k_1}(A_1))^T \text{best}_{k_2}(A_2)\|, \|(A_2 - \text{best}_{k_2}(A_2))^T \text{best}_{k_1}(A_1)\|\}.\end{aligned}$$

REMARK. It is easy to see that each of the corollaries following Theorem 2.1 are direct consequences of the corresponding corollaries established above.

Perturbation Results for Matrices with Low-Rank-Plus-Shift Structure. Now we return to matrices with low-rank-plus-shift structure, and we consider the case the structure is only approximately satisfied. It turns out that the way this

approximation is quantified has direct impact on the perturbation bounds we can derive. In the following we provide two theorems one with $O(\sqrt{\epsilon})$ and the other with $O(\epsilon)$. The difference in the assumptions for the derivation of these two results is very subtle, but it gives rise to qualitatively different results.

To derive the perturbation bounds, we first need two technical lemmas which were proved in [12].

LEMMA 3.7. *Assume the following equality*

$$\begin{bmatrix} A & B^T \\ B & C \end{bmatrix} = X + E$$

holds for some symmetric positive semi-definite matrix X Then we have

$$\|B\| \leq \sqrt{(\|A\| + \|E\|)(\|C\| + \|E\|)}.$$

LEMMA 3.8. *Let the symmetric matrix Z be partitioned as*

$$Z = \begin{bmatrix} A & B^T \\ B & C \end{bmatrix}.$$

Then $\|Z\| \leq \max\{\|A\|, \|C\|\} + \|B\|$.

THEOREM 3.9. *Let $A = [A_1, A_2] \in \mathcal{R}^{m\times n}$. Assume that for some integer $k < \min\{m, n\}$ there exists $\epsilon \geq 0$ for which the eigenvalues of $X \equiv A^T A - \sigma^2 I$ satisfy*[2]

$$\begin{aligned} \lambda_j(X) &> 3\epsilon + \eta, \quad j \leq k, \\ |\lambda_j(X)| &\leq \epsilon, \qquad j > k, \end{aligned}$$

where $\eta = 2\sqrt{\|X\|\epsilon + \epsilon^2} = O(\sqrt{\epsilon})$. Partiton $X = (X_{ij})_{i,j=1}^2$ conformally with that of A. Define k_i such that

$$\begin{aligned} \lambda_j(X_{ii}) &> \epsilon, \quad j \leq k_i, \\ |\lambda_j(X_{ii})| &\leq \epsilon, \quad j > k_i, \end{aligned}$$

for $i = 1, 2$. Then

$$\|\text{best}_k(A) - \text{best}_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_1)])\| \leq \frac{\|A\|\eta}{\sigma_k^2(A) - \sigma_{k+1}^2(A) - \eta}.$$

Proof. By the eigendecomposition of A and the assumptions of its eigenvalues, we can write $X = Y + E$, where $Y^T E = 0$, and Y is positive semi-definite with $\text{rank}(Y) = k$, $\|E\| \leq \epsilon$, and

$$\lambda_k(Y) = \lambda_k(X) > 3\epsilon + \eta.$$

On the other hand, using the partition of A, we can write

$$X = \begin{bmatrix} A_1^T A_1 - \sigma^2 I & A_1^T A_2 \\ A_2^T A_1 & A_2^T A_2 - \sigma^2 I \end{bmatrix}.$$

[2] We assume that the eigenvalues of a matrix X are ordered in nonincreasing order $\lambda_1(X), \geq \ldots, \lambda_n(X)$.

Now for $i = 1, 2$, write the SVD fo A_i as follows

$$A_i = [U_{i1}, U_{i2}]\,\mathrm{diag}(\Sigma_{i1}, \Sigma_{i1})[V_{i1}, V_{i2}]^T,$$

where $\Sigma_{i1} = \mathrm{diag}(\sigma_{i1}, \ldots, \sigma_{i,k_i})$ and $\Sigma_{i2} = \mathrm{diag}(\sigma_{i,k_i+1}, \ldots, \sigma_{i,m_i})$. By definition the integers k_i are chosen such that

$$\begin{aligned} \sigma_{ij}^2 - \sigma^2 > \epsilon \quad & j \le k_i, \\ |\sigma_{ij}^2 - \sigma^2| \le \epsilon \quad & j > k_i, \end{aligned}$$

for $i = 1, 2$. Or equivalently, $\lambda_j(A_i^T A_i - \sigma^2 I) > \epsilon$ for $j \le k_i$, and $|\lambda_j(A_i^T A_i - \sigma^2 I)| \le \epsilon$ for $j > k_i$. It is easy to see that $k_i \le k$ since $\sigma_{ij} \le \sigma_j(A)$.

Next we write $A = BW^T \equiv [B_1, B_2]W^T$, where

$$B_1 = [U_{11}\Sigma_{11}, U_{21}\Sigma_{21}], \quad B_2 = [U_{12}\Sigma_{12}, U_{22}\Sigma_{22}],$$

and

$$W = \begin{bmatrix} V_{11} & 0 & V_{12} & 0 \\ 0 & V_{21} & 0 & V_{22} \end{bmatrix}.$$

Without loss of generality, we assume that W is orthogonal. (Otherwise replace W and B_2 by $[W, W^\perp]$ and $[B_2, 0]$, respectively.) Define

$$\Delta = \mathrm{best}_k(A) - \mathrm{best}_k([\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_1)]).$$

It can be verified that

$$\mathrm{best}_k[\mathrm{best}_{k_1}(A_1), \mathrm{best}_{k_2}(A_2)] = \mathrm{best}_k[B_1, 0]W^T, \quad \mathrm{best}_k(A) = \mathrm{best}_k(B)W^T,$$
$$\|\Delta\| = \|\mathrm{best}_k(B) - \mathrm{best}_k([B_1, 0])\|.$$

Now in order to apply Corollary 3.4, we need to verify the condition $\sigma_k(B_1) > \sigma_{k+1}(B)$, and derive a lower bound on $\sigma_k(B_1)^2 - \sigma_{k+1}^2(B)$ and an upper bound on $\|B_2^T B_1\|$. (Notice that $\|B_2^T \mathrm{best}_k(B_1)\| \le \|B_2^T B_1\|$.) To this end,

1) we apply Lemma 3.7 twice to obtain an upper bound on $\|B_2^T B_1\|$. It is easy to see that both $B^T B - \sigma^2 I$ and $B_2^T B_2 - \sigma^2 I$ can be written as the sum of a symmetric positive semi-definite matrix and a matrix with norm no greater than ϵ. Applying Lemma 3.7 to

$$B^T B - \sigma^2 I = \begin{bmatrix} B_1^T B_1 - \sigma^2 I & B_1^T B_2 \\ B_2^T B_1 & B_2^T B_2 - \sigma^2 I \end{bmatrix}$$

gives

$$\begin{aligned} \|B_2^T B_1\| &\le \sqrt{(\|B_1^T B_1 - \sigma^2 I\| + \epsilon)(\|B_2^T B_2 - \sigma^2 I\|\epsilon)} \\ &\le \sqrt{(\|X\| + \epsilon)(\|B_2^T B_2 - \sigma^2 I\|\epsilon)}. \end{aligned}$$

Apply Lemma 3.7 to

$$B_2^T B_2 - \sigma^2 I = \begin{bmatrix} \Sigma_{12}^2 - \sigma^2 I & \Sigma_{12} U_{12}^T U_{22} \Sigma_{22} \\ \Sigma_{22} U_{22}^T U_{12} \Sigma_{12} & \Sigma_{22}^2 - \sigma^2 I \end{bmatrix}$$

yields

$$\|\Sigma_{12} U_{12}^T U_{22} \Sigma_{22}\| \le (\|\Sigma_{12}^2 - \sigma^2 I\| + \epsilon)(\|\Sigma_{22}^2 - \sigma^2 I\| + \epsilon) \le 4\epsilon^2,$$

where we have used $\|\Sigma_{i2}^2 - \sigma^2 I\| \le \epsilon$. By Lemma 3.8, we obtain $\|B_2^T B_2 - \sigma^2 I\| \le 3\epsilon$ and hence

$$\|B_2^T B_1\| \le 2\sqrt{\|X\|\epsilon + \epsilon^2} \equiv \eta.$$

2) we now give an lower bound on $\sigma_k(B_1)^2 - \sigma_{k+1}^2(B)$. Write

$$B^T B - \sigma^2 I = \begin{bmatrix} B_1^T B_1 - \sigma^2 I & \\ & B_2^T B_2 - \sigma^2 I \end{bmatrix} + \begin{bmatrix} & B_1^T B_2 \\ B_2^T B_1 & \end{bmatrix}$$

Use perturbation bounds for eigenvalues, we have

$$\lambda_k(X) = \lambda_k(B^T B - \sigma^2 I) \le \lambda_k(\operatorname{diag}(B_1^T B_1 - \sigma^2 I, B_2^T B_2 - \sigma^2 I)) + \eta.$$

The condition $\lambda_k(X) > 3\epsilon + \eta$ implies

$$\lambda_k(B_1^T B_1 - \sigma^2 I) > \|B_2^T B_2 - \sigma^2 I\|,$$

because $\|B_2^T B_2 - \sigma^2 I\| \le 3\epsilon$. Thus

$$|\sigma_k^2(B_1) - \sigma_k^2(A)| = |\lambda_k(B_1^T B_1 - \sigma^2 I) - \lambda_k(X)| \le \|B_1^T B_2\| \le \eta.$$

It follws that

$$\sigma_k^2(B_1) - \sigma_{k+1}^2(B) \ge \sigma_k^2(A) - \eta - \sigma_{k+1}^2(A) \ge \epsilon + \eta > 0.$$

Finally, by Corollary 3.4, we have

$$\|\Delta\| \le \frac{\|B\|\|B_1^T B_2\|}{\sigma_k^2(B_1) - \sigma_{k+1}^2(B)} \le \frac{\|A\|\eta}{\sigma_k^2(A) - \sigma_{k+1}^2(A) - \eta},$$

completing the proof. □

EXAMPLE 1. Let s be small, and for any $\sigma > s$, define

$$c_1 = \sqrt{1 + \sigma^2 + s^2}, \quad c_1 = \sqrt{\sigma^2 + s^2}, \quad c_3 = \sqrt{\sigma^2 + s^2}, \quad \epsilon = c_1^2 s^2.$$

Let $A = [A_1, A_2]$ with

$$A_1 = \frac{1}{\sqrt{1+s^2}} \begin{bmatrix} D \\ sDJ \end{bmatrix}, \quad A_2 = \frac{1}{\sqrt{1+s^2}} \begin{bmatrix} -sDJ \\ D \end{bmatrix}, \quad \text{with } J = \begin{bmatrix} & & 1 \\ & 1 & \\ 1 & & \end{bmatrix},$$

and $D = \operatorname{diag}(c_1, c_2, c_3)$. It follows that

$$A = \frac{1}{\sqrt{1+s^2}} \begin{bmatrix} D & \\ & D \end{bmatrix} \begin{bmatrix} I & -sJ \\ sJ & I \end{bmatrix}.$$

It can be verified that

$$\begin{aligned} &\lambda_{1,2}(A^T A - \sigma^2 I) = 1 + s^2 > \epsilon > |\lambda_{1,2}(A^T A - \sigma^2 I)|, && j \ge 3 \\ &\lambda_1(A_i^T A_i - \sigma^2 I) = 1 + s^2(1 + \sigma^2 - s^2) > \epsilon \ge |\lambda_{1,2}(A^T A - \sigma^2 I)|, && j \ge 2, \end{aligned}$$

for $i = 1, 2$. Hence $k = 2$, and $k_1 = k_2 = 1$. It can be verified that

$$\operatorname{best}_k(A) = \frac{c_1}{\sqrt{1+s^2}}[e_1, e_3][e_1 - se_6, se_3 + e_4]^T,$$

$$\text{best}_{k_1}(A_1) = \frac{1}{\sqrt{1+s^2}}[ce_1 + sc_3e_6, 0, 0], \quad \text{best}_{k_1}(A_1) = \frac{1}{\sqrt{1+s^2}}[-sc_3e_3 + c_1e_4, 0, 0].$$

So we have

$$\text{best}_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]) = [\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)].$$

and

$$\|\text{best}_k(A) - [\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]\| = \frac{sc_1}{\sqrt{1+s^2}} = \sqrt{\frac{\epsilon}{1+s^2}}.$$

THEOREM 3.10. *Let $A = [A_1, A_2]$. If there exists $\epsilon < \sigma^2$ and integer k such that*

$$\lambda_k(A^TA - \sigma^2 I) > \epsilon \geq |\lambda_j(A^TA - \sigma^2 I)|,$$

for $j \geq k+1$, and $\lambda_k(A_i^TA_i - \sigma^2 I) > \epsilon, i = 1, 2$. Then

$$\|(A_1 - \text{best}_k(A_1))^TA_2\| \leq \eta_1, \quad \|(A_2 - \text{best}_k(A_2))^TA_1\| \leq \eta_2$$

and

$$\|\text{best}_k(A) - \text{best}_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)])\| \leq \frac{\|A\|\eta}{\lambda_{\max} - \epsilon}, \tag{3.5}$$

where

$$\eta_i = \left(\sigma + 2\|A\| + \frac{2\|A\|^3}{\lambda_k(A_i^TA_i - \sigma^2 I)}\right)\frac{\epsilon}{\sigma + \sqrt{\sigma^2 - \epsilon}}, \quad i = 1, 2,$$

$$\eta = \max\{\eta_1, \eta_2\} = \left(\sigma + 2\|A\| + \frac{2\|A\|^3}{\lambda_{\min}}\right)\frac{\epsilon}{\sigma + \sqrt{\sigma^2 - \epsilon}},$$

$$\lambda_{\max} = \max_{i=1,2}\lambda_k(A_i^TA_i - \sigma^2 I), \quad \lambda_{\min} = \min_{i=1,2}\lambda_k(A_i^TA_i - \sigma^2 I).$$

Proof. Denote $\lambda_j = \lambda_j(A^TA - \sigma^2 I)$, and

$$\Lambda_1 = \text{diag}(\lambda_1, \ldots, \lambda_k), \quad \Lambda_2 = \text{diag}(\lambda_{k+1}, \ldots, \lambda_n).$$

The the eigendecomposition of $A^TA - \sigma^2 I$ and the SVD of A can be written as

$$A^TA - \sigma^2 I = V\,\text{diag}(\Lambda_1, \Lambda_2)V^T, \quad A = U\,\text{diag}(\sqrt{\Lambda_1^2 + \sigma^2 I}, \sqrt{\Lambda_2^2 + \sigma^2 I})V^T$$

for some orthogonal matrices U and V. Let $E = U\,\text{diag}(0, \sigma I - \sqrt{\Lambda_2^2 + \sigma^2 I})$. It can be verified that $\|E\| \leq \epsilon/(\sigma + \sqrt{\sigma^2 - \epsilon}) \equiv \tau$, and

$$\tilde{A} \equiv A + E = U\,\text{diag}(\sqrt{\Lambda_1^2 + \sigma^2 I}, \sigma I)V^T$$

has the low-rank-plus-shift structure. Now partition

$$E = [E_1, E_2], \quad \tilde{A} = [\tilde{A}_1, \tilde{A}_2]$$

conformally as that of A. Then $\|E_i\| \le \tau$. Since

$$A^TA - \sigma^2 I = \tilde{A}^T\tilde{A} - \sigma^2 I + \tilde{E}, \quad \tilde{E} = V\,\mathrm{diag}(0, \Lambda_2)V^T = (\tilde{E}_{i,j})_{i,j=1}^2.$$

It can be verified that $A_i^TA_i - \sigma^2 I = \tilde{A}_i^T\tilde{A}_i - \sigma^2 I + \tilde{E}_{ii}$, and we have

$$\begin{aligned}\lambda_k(\tilde{A}_i^T\tilde{A}_i - \sigma^2 I) &\ge \lambda_k(A_i^TA_i - \sigma^2 I) - \|\tilde{E}_{ii}\| \\ &\ge \lambda_k(\tilde{A}_i^T\tilde{A}_i - \sigma^2 I) - \epsilon > 0\end{aligned}$$

It follows that $\mathrm{rank}(\tilde{A}_i^T\tilde{A}_i - \sigma^2 I) = k$. By Theorem 2.7, we have

$$(\tilde{A}_1 - \mathrm{best}_k(\tilde{A}_1))^T\tilde{A}_2 = 0, \quad (\tilde{A}_2 - \mathrm{best}_k(\tilde{A}_2))^T\tilde{A}_1 = 0.$$

Denote $\Delta_i = \mathrm{best}_k(\tilde{A}_i) - \mathrm{best}_k(A_i) - E_i$. Then

$$\|\Delta_i\| \le \|\mathrm{best}_k(\tilde{A}_i) - \mathrm{best}_k(A_i)\| + \|E_i\|$$

It follows from Corollary 3.4 that

$$\|\mathrm{best}_k(\tilde{A}_i) - \mathrm{best}_k(A_i)\| \le \left(1 + \frac{\|A\|^2}{\sigma_k^2(A_i) - \sigma_{k+1}^2(\tilde{A}_i)}\right)\|E_i\| = \left(1 + \frac{\|A\|^2}{\sigma_k^2(A_i) - \sigma^2}\right)\|E_i\|$$

and therefore

$$\|\Delta_i\| \le \left(2 + \frac{2\|A\|^2}{\lambda_k(A_i^TA_i)}\right)\tau,$$

here we have used $\|A_i\| \le \|A\|$, $\|\tilde{A}_i\| \le \|\tilde{A}\|$, and $\sigma_{k+1}(\tilde{A}_i) = \sigma$. Since

$$A_i - \mathrm{best}_k(A_i) = \tilde{A}_i - \mathrm{best}_k(\tilde{A}_i) + \Delta_i,$$

we have

$$\begin{aligned}\|(A_1 - \mathrm{best}_k(A_1))^TA_2\| &= \|\Delta_1^TA_2 - (\tilde{A}_1 - \mathrm{best}_k(\tilde{A}_1)\| \\ &\le \|A_2\|\|\Delta_1\| + \sigma_{k+1}(\tilde{A}_1)\|E_2\| \\ &\le (\sigma + 2\|A\| + 2\|A\|^3/\lambda_k(A_1^TA_1 - \sigma^2 I))\tau = \eta_1.\end{aligned}$$

We can similarly prove $\|(A_2 - \mathrm{best}_k(A_2))^TA_1\| \le \eta_2$. Now it follows from

$$\sigma_k([\mathrm{best}_k(A_1), \mathrm{best}_k(A_2)]) \ge \max_{i=1,2}\sigma_k(A_i),$$

that

$$\sigma_k^2([\mathrm{best}_k(A_1), \mathrm{best}_k(A_2)]) - \sigma_{k+1}^2(A) \ge \max_{i=1,2}\lambda_k(A_i^TA_i - \sigma^2 I) - \epsilon = \lambda_{\max} - \epsilon.$$

Finally Corollary 2.4 and the above give (3.5). □

REMARK. We notice that in order for the perturbation bound to be of order $O(\epsilon)$, $\lambda_{\min}$ needs to be of order $O(1)$.

EXAMPLE 2. Now we construct a class of matrices that satisfy the conditions of Theorem 3.10. For any orthonormal matrices U_1 and V_1 with k columns, let $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_k)$, where $\lambda_i \gg \sigma^2 > 0$ for $i = 1, \ldots, k$. Let

$$D_1 = (\Lambda\sqrt{\Lambda^2 + \epsilon I})^{-1}(\Lambda^2 - (\sigma^2 - \epsilon)I), \quad D_2 = \sqrt{I - D_1^2}, \quad U_2 = [U_1, U_1^{\perp}][D_1, D_2]^T;$$

where $U_1^{\perp}$ is any orthonormal matrix of k columns that is orthogonal to U_1. Define

$$A_1 = U_1 \Lambda V_1^T, \quad A_2 = U_2 \sqrt{\Lambda^2 + \epsilon I} V_1^T.$$

It follows that

$$X = [A_1, A_2]^T [A_1, A_2] - \sigma^2 I = \begin{bmatrix} A_1^T A_1 - \sigma^2 I & A_1^T A_2 \\ A_2^T A_1 & A_2^T A_2 - \sigma^2 I \end{bmatrix}$$

$$= \begin{bmatrix} A_1^T A_1 - \sigma^2 I & A_1^T A_1 - \sigma^2 I + \epsilon \\ A_2^T A_1 & A_2^T A_1 \end{bmatrix} = \begin{bmatrix} A_2^T A_2 - \sigma^2 I & A_2^T A_2 - \sigma^2 I \\ A_2^T A_2 - \sigma^2 I & A_2^T A_2 - \sigma^2 I \end{bmatrix} - \begin{bmatrix} \epsilon & 0 \\ 0 & 0 \end{bmatrix}.$$

Hence,

$$\begin{aligned} &\lambda_j(X) \geq 2(\lambda_j - \sigma^2 + \epsilon) \gg \epsilon, && j \leq k, \\ &|\lambda_j(X)| \leq \epsilon, && j > k, \end{aligned}$$

and by definition $k_1 = k_2 = k$.

REMARK. For the case where $k_1 < k$ and $k_2 < k$, if we replace $\text{best}_{k_1}(A_1)$ and $\text{best}_{k_2}(A_2)$ by $\text{best}_k(A_1)$ and $\text{best}_k(A_2)$, the error

$$||\Delta|| = ||\text{best}_k(A) - \text{best}_k[\text{best}_k(A_1), \text{best}_k(A_2)]||$$

may still be $O(\sqrt{\epsilon})$. For example, in Example 1, we have

$$\text{best}_2[\text{best}_2(A_1), \text{best}_2(A_2)] = \text{best}_2[\text{best}_1(A_1), \text{best}_1(A_2)].$$

Therefore,

$$||\text{best}_2(A) - \text{best}_2[\text{best}_2(A_1), \text{best}_2(A_2)]|| = \sqrt{\frac{\epsilon}{1+s^2}}.$$

REFERENCES

[1] M. Berry. Large Scale Singular Value Computations. *International Journal of Supercomputer Applications*, 6:13-49, 1992.

[2] J. Cullum, R. A. Willoughby and M. Lake. A Lanczos Algorithm for Computing Singular Values and Vectors of Large Matrices. *SIAM Journal on Scientific and Statistical Computing*, 4:197–215, 1983.

[3] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, Maryland, 2nd edition, 1989.

[4] B.N. Parlett. *The Symmetric Eigenvalue Problem*. SIAM Press, Philadelphia, 1998.

[5] H. Simon and H. Zha. Low Rank Matrix Approximation Using the Lanczos Bidiagonalization Process. CSE Tech. Report CSE-97-008, 1997. (Also LBNL Tech. Report LBNL-40767-UC-405.) Submitted to *SIAM Journal on Scientific Computing*

[6] G. Xu and T. Kailath. Fast subspace decompsotion. *IEEE Transactions on Signal Processing*, 42:539–551, 1994.

[7] G. Xu, H. Zha, G. Golub, and T. Kailath. Fast algorithms for updating signal subspaces. *IEEE Transactions on Circuits and Systems*, 41:537–549, 1994.

[8] H. Zha and H.D. Simon. On Updating Problems in Latent Semantic Indexing. Technical Report CSE-97-008, Department of Computer Science and Engineering, The Pennsylvania State University, 1997. To appear in *SIAM Journal on Scientific Computing*.

[9] H. Zha and H. Simon. A Subspace-based model for Latent Semantic Indexing in information retrieval. To appear in *Proceedings of Interface '98*.

[10] H. Zha, O. Marques and H. Simon. Large-scale SVD and Subspace-based Methods for Information Retrieval. *Proceedings of Irregular '98*, Lecture Notes in Computer Science, 1457:29–42, 1998.

[11] H. Zha and Z. Zhang. A sorted partial Jacobi method and its convergence analysis. *Linear Algebra and its Applications*, 270:79-108, 1997.

[12] H. Zha and Z. Zhang. Matrices with low-rank-plus-shift structure: partial SVD and latent semantic indexing. To appear in *SIAM Journal on Matrix Analysis and Applications*, 1998.

ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY
ONE CYCLOTRON ROAD, BERKELEY, CALIFORNIA 94720

Prepared for the U.S. Department of Energy under Contract No. DE-AC03-76SF00098