

Lawrence Berkeley National Laboratory

Recent Work

Title

On matrices with low-rank-plus-shift structure: partial SVD and latent semantic indexing

Permalink

<https://escholarship.org/uc/item/8093v13x>

Author

Zha, Hongyuan

Publication Date

1998-08-01




ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY

On Matrices with Low-Rank-Plus-Shift Structure: Partial SVD and Latent Semantic Indexing

Hongyuan Zha and Zhenyue Zhang

**Computing Sciences Directorate
National Energy Research
Scientific Computing Division**

August 1998



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

**On Matrices with Low-Rank-Plus-Shift Structure:
Partial SVD and Latent Semantic Indexing**

Hongyuan Zha and Zhenyue Zhang

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

August 1998

This work was supported in part by the Director, Office of Energy Research, Office of Laboratory Policy and Infrastructure Management, Office of Computational and Technology Research, Division of Mathematical, Information, and Computational Sciences, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098, by National Science Foundation Grant CCR-9619452, and by NSFC Project No. 19771073. Worked performed at the Ernest Orlando Lawrence Berkeley National Laboratory.

ON MATRICES WITH LOW-RANK-PLUS-SHIFT STRUCTURE: PARTIAL SVD AND LATENT SEMANTIC INDEXING

HONGYUAN ZHA* AND ZHENYUE ZHANG†

Abstract. We present a detailed analysis of matrices satisfying the so-called *low-rank-plus-shift* property in connection with the computation of their partial singular value decomposition. The application we have in mind is Latent Semantic Indexing for information retrieval where the term-document matrices generated from a text corpus approximately satisfy this property. The analysis is motivated by developing more efficient methods for computing and updating partial SVD of large term-document matrices and gaining deeper understanding of the behavior of the methods in the presence of noise.

1. Introduction. In many applications such as compression of multiple-spectral image cubes, regularization methods for ill-posed problems, latent semantic indexing in information retrieval for large document collections, it is necessary to find a low rank approximation of a given large and/or sparse matrix $A \in \mathcal{R}^{m \times n}$ [11]. The theory of singular value decomposition (SVD) provides the following characterization of the best low rank approximation of A in terms of Frobenius norm $\|\cdot\|_F$ [6].

THEOREM 1.1. *Let the singular value decomposition of $A \in \mathcal{R}^{m \times n}$ be $A = P\Sigma Q^T$ with $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_{\min(m,n)})$, $\sigma_1 \geq \dots \geq \sigma_{\min(m,n)}$, and P and Q orthogonal. Then for $1 \leq j \leq \min(m, n)$,*

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

And the minimum is achieved with $A_j \equiv P_j \text{diag}(\sigma_1, \dots, \sigma_j) Q_j^T$, where P_j and Q_j are the matrices formed by the first j columns of P and Q , respectively.

It follows from Theorem 1.1 that once the SVD of A is available, the best rank- j approximation of A is readily computed. We call $A_j = P_j \text{diag}(\sigma_1, \dots, \sigma_j) Q_j^T$ a partial SVD of A . The state-of-the-art methods for computing the partial SVD of large and/or sparse matrices are based on variants of Lanczos algorithms and the core computation at each iterative steps involves matrix-vector multiplications [9]. In order to effectively deal with large-scale problems, one is required to exploit various structures of the matrices. Despite its importance, the exploitation of structures so far has been restricted to 1) using the sparsity of a sparse matrix, 2) using displacement-rank structures such as Toeplitz or Hankel structure of the matrix, to accelerate the matrix-vector multiplications used in the Lanczos process. In this paper, however, we propose to explore an alternative structure that is based on the singular value

* 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 the Director, Office of Energy Research, Office of Laboratory Policy and Infrastructure Management, Office of Computational and Technology Research, Division of Mathematical, Information, and Computational Sciences, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098, and by NSF grant CCR-9619452.

† 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, and Scientific Research Foundation for Returned Overseas Chinese Scholars, State Education Commission.

spectrum of a matrix. Specifically, we investigate matrices possessing the so-called *low-rank-plus-shift* structure, i.e., those matrices A (approximately) satisfying

$$(1.1) \quad A^T A = \text{a low rank matrix} + \text{a multiple of the identity matrix.}$$

The particular application we have in mind is Latent Semantic Indexing (LSI) for information retrieval and we will show by way of examples that the term-document matrices generated from text corpora approximately satisfy Equation (1.1). In large-scale LSI applications such as the World Wide Web, the term-document matrix generated is usually very large and can not be kept in RAM or disk. In [15, 16] we have shown that the low-rank-plus-shift structure of the term-document matrix A allows us to compute its partial SVD in a block-wise fashion whereby partial SVD of submatrices of A are computed separately and then merged to obtain a partial SVD of A . The purpose of this paper is to further analyze the properties of matrices with low-rank-plus-shift structure especially when Equation (1.1) is only approximately satisfied. We hope our analysis will provide deeper insights into this special class of matrices which will enable us to develop more efficient methods for computing their partial SVD.

The rest of the paper is organized as follows: In Section 2, we provide a brief background on LSI and review some of the results in [14, 15, 16] related to computing the partial SVD of term-document matrices. In Section 3, we discuss some matrix approximation problems associated with the low-rank-plus-shift structure, and show by way of examples that the term-document matrices generated from text corpora approximately satisfy Equation (1.1). In Section 4 we prove a result on the partial SVD of a block-column partitioned matrix with low-rank-plus-shift structure. This result enables us to efficiently compute the partial SVD either with an incremental approach or divide-and-conquer approach. We will also discuss the ramification of the result in dealing with LSI updating problems. In Section 5 we provide a perturbation analysis of the result when the low-rank-plus-shift property is only satisfied approximately. In Section 6 we conclude the paper with some remarks on future research.

2. Latent Semantic Indexing. Latent semantic indexing is a concept-based automatic indexing method that aims at overcoming the two fundamental problems which plague traditional lexical-matching indexing schemes: synonymy and polysemy [2, 5]. Synonymy refers to the problem that several different words can be used to express a concept and the keywords in a user's query may not match those in the relevant documents while polysemy means that words can have multiple meanings and user's words may match those in irrelevant documents [8]. LSI is an extension of the vector space model for information retrieval [7, 10]. In the vector space model, the collection of text documents is represented by a *term-document* matrix $A = [a_{ij}] \in \mathcal{R}^{m \times n}$, where a_{ij} is the number of times term i appears in document j , and m is the number of terms and n is the number of documents in the collection. Consequently, a document becomes a column vector, and a user's query can also be represented as a vector of the same dimension. The similarity between a query vector and a document vector is usually measured by the cosine of the angle between them, and for each query a list of documents ranked in decreasing order of similarity is returned to the user. LSI extends this vector space model by modeling the term-document relationship using the singular value decomposition (SVD) of the term-document matrix A . Specifically, using the notation in Theorem 1.1, we substitute A by its best rank- k approximation $A_k \equiv P_k \Sigma_k Q_k^T$, where Σ_k is the k -th leading principal submatrix of Σ . Corresponding

to each of the k reduced dimensions is associated a latent concept which may not have any explicit semantic content yet helps to discriminate documents [2, 5].

Large text corpora such as those generated from World Wide Web give rise to very large term-document matrices, and the computation of their partial SVD poses a very challenging problem. Fortunately the term-document matrices possess certain useful properties besides sparsity that can be exploited for this matter. In [15, 16], we developed a theoretical foundation for LSI using the concept of subspaces, and we showed that the model we proposed imposes a so-called *low-rank-plus-shift* structure that is approximately satisfied by the cross-product of the term-document matrices.¹ Specifically, we showed that the term-document matrix $A \in \mathcal{R}^{m \times n}$ satisfies

$$(2.2) \quad A^T A/m \approx CWC^T + \sigma^2 I_n,$$

where $C \in \mathcal{R}^{n \times k}$ is the matrix whose columns represent latent concepts, $W \in \mathcal{R}^{k \times k}$ is a symmetric positive definite matrix, and σ is the variance of the noise. In LSI applications $k \ll \min\{m, n\}$, justifying the use of the terminology *low-rank-plus-shift* structure.

In [14], we considered the updating problems for LSI: Let A be the term-document matrix for the original text collection and D represents a collection of new documents. The goal is to compute the partial SVD of $[A, D]$. However, in LSI applications, only A_k for some chosen k is available and the matrix A has been discarded. Since updating in this situation is based on a low-rank approximation of A , it has been argued in the literature that one will not be able to get an accurate partial SVD of $[A, D]$. In Section 4, we show, however, that this is not the case since $[A, D]$ has the low-rank-plus-shift structure [14]. We will show that no retrieval accuracy degradation will occur if updating is done with a proper implementation. In [15, 16], we also discussed how to compute the partial SVD of a term-document matrix in a block-column partitioned form $A = [A_1, A_2]$ using a divide-and-conquer approach whereby the partial SVDs of A_1 and A_2 are first computed and the results are then merged into a partial SVD of A . This approach is rich in coarse-grain parallelism and can be used to handle very large term-document matrices. The justification for this divide-and-conquer approach will be discussed in greater detail in Section 4, and perturbation analysis will be provided to show that the approach is still valid even if the term-document matrix A only approximately satisfies the low-rank-plus-shift structure.

3. A Matrix Approximation Problem. From the discussion in Section 2 we know that the term-document matrix A approximately satisfies the low-rank-plus-shift property and therefore A should have flat trailing singular values. In this section we use several example text collections to illustrate this issue. In order to assess whether a given matrix has the low-rank-plus-shift property, we investigate the following matrix approximation problem: Given a general rectangular matrix, what is the *closest* matrix that has the low-rank-plus-shift property. To proceed we first define a matrix set for a given $k > 0$,

$$\mathcal{J}_k = \{B \in \mathcal{R}^{m \times n} \mid \sigma_1(B) \geq \dots \geq \sigma_{\min\{m, n\}}(B), \sigma_{k+1}(B) = \dots = \sigma_{\min\{m, n\}}(B)\}.$$

With this notation, the matrix approximation problem reduces to finding the distance between a general matrix A and the set \mathcal{J}_k . In the following we consider the cases where distance is defined either by Frobenius norm $\|\cdot\|_F$ or spectral norm $\|\cdot\|_2$.

¹ The low-rank-plus-shift structure was first discussed in the context of array signal processing [12, 13, 17].

THEOREM 3.1. *Let the SVD of A be $A = U\Sigma V^T$, $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_{\min\{m,n\}})$, and U and V orthogonal. Then for $k \leq \min\{m, n\}$,*

$$(3.3) \quad \arg \min_{J \in \mathcal{J}_k} \|A - J\|_p = U_k \Sigma_k V_k^T + \tau_p U_k^\perp (V_k^\perp)^T,$$

where $\Sigma_k = \text{diag}(\sigma_1, \dots, \sigma_k)$, $U = [U_k, U_k^\perp]$ and $V = [V_k, V_k^\perp]$. Furthermore,

$$\tau_p = \begin{cases} \sum_{i=k+1}^{\min\{m,n\}} \sigma_i / (\min\{m, n\} - k), & p = F \\ (\sigma_{k+1} + \sigma_{\min\{m,n\}}) / 2, & p = 2. \end{cases}$$

Proof. First define

$$\min_p = \begin{cases} \|a_k - \tau_p\|, & p = F \\ \|a_k - \tau_p\|_\infty, & p = 2. \end{cases}$$

where $a_k = [\sigma_{k+1}, \dots, \sigma_{\min\{m,n\}}]$. It is readily checked that \min_p is achieved by the right-hand side of Equation (3.3). Therefore all we need to prove is $\|A - J\|_p \geq \min_p$ for $p = 2, F$ and for any $J \in \mathcal{J}_k$.

To this end we use standard perturbation analysis of singular values which states that [6]

$$|\sigma_i(A) - \sigma_i(J)| \leq \|A - J\|_2, \quad i = 1, \dots, \min\{m, n\},$$

and

$$\sum_{i=1}^{\min\{m,n\}} (\sigma_i(A) - \sigma_i(J))^2 \leq \|A - J\|_F^2.$$

It follows that

$$\max_{k+1 \leq i \leq \min\{m,n\}} |\sigma_i(A) - \sigma_i(J)| \leq \max_{1 \leq i \leq \min\{m,n\}} |\sigma_i(A) - \sigma_i(J)| \leq \|A - J\|_2,$$

and

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

Notice that $\sigma_{k+1}(J) = \dots = \sigma_{\min\{m,n\}}(J)$, it can be readily verified that the minima of the left-hand sides of the above two inequalities, i.e., \min_p , $p = 2, F$, are achieved by τ_2 and τ_F , respectively. \square

EXAMPLES. In the following we will apply the above theorem to two example text collections and see how close the associated term-document matrices are to the set of matrices with low-rank-plus-shift structure. Our first example is the MEDLINE collection from the Cornell SMART system [3]. The term-document matrix is of size 3681×1033 . The singular value distribution is plotted on the left of Figure 1. Our second example is from a collection consisting of news articles from 20 newsgroups [4]. The term-document matrix is of size 33583×1997 . Its singular values are plotted on the right of Figure 1. From Theorem 3.1, the best approximation from \mathcal{J}_k to A is given by

$$A^{(k)} \equiv U_k \Sigma_k V_k^T + \tau_p U_k^\perp (V_k^\perp)^T$$




FIG. 1. *Singular value distributions: 3681 × 1033 term-document matrix of MEDLINE Collection (left) and 33583 × 1997 NEWSGROUP Collection (right)*

For the MEDLINE collection we have $\|A - A^{(100)}\|_F / \|A\|_F = 0.2909$ and for the NEWSGROUP collection we have $\|A - A^{(100)}\|_2 / \|A\|_2 = 0.0491$. Several other text collections from the Cornell SMART system have also been tested and we observed similar singular value distributions: initially the singular values decrease rapidly and then the spectrum curve levels off, but the singular values are never close to zero. (Unless the sparse term-document matrix is *structurally* rank-deficient.) The last point is very important: we usually should not treat those matrices simply as near rank-deficient and it is more appropriate that the more general low-rank-plus-shift structure with a nonzero σ be used (cf. Equation (2.2)).

4. The Low-Rank-Plus-shift Structure. We start with an examination of the changes of the singular values of a matrix when its elements undergo certain type of modifications. If some of the elements of a general matrix is set to zero, generally it is not possible to tell whether the singular values of the matrix will increase or decrease. However, a result we will show below states that the singular values of a matrix will always decrease if some submatrices of the matrix are replaced by its low-rank approximations. To proceed we introduce some notation: for any matrix $A \in \mathcal{R}^{m \times n}$, we will use $\text{best}_k(A)$ to denote its best rank- k approximation (cf. Theorem 1.1), and its singular values are assumed to be arranged in nonincreasing order,

$$\sigma_1(A) \geq \sigma_2(A) \geq \dots \geq \sigma_m(A).$$

As a convention when we compare the singular values of two matrices with the same number of rows but different number of columns we will count the singular values according to the number of rows. With the above preparation we present our first result. The proof is similar to that of a slightly special case presented in [14] and therefore is omitted.

THEOREM 4.1. *Let $A \in \mathcal{R}^{m \times n}$ and write $A = [A_1, A_2]$. Then for any k_1 and k_2 , we have*

$$\sigma_i([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]) \leq \sigma_i([A_1, A_2]), \quad i = 1, \dots, m.$$

REMARK. It is not true that replacing *arbitrary* submatrices of a matrix by their low-rank approximations will result in the decrease of its singular values as is

illustrated in the following example: Let

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

Notice that \hat{A} is obtained from A by replacing its (1, 1) and (2, 2) elements by zero (a best rank-zero approximation). Even though the largest singular value decreases, its smallest singular value increases.

It is rather easy to find examples for which strict *inequalities* hold in Theorem 4.1. In the following we show that this will not be the case if A has the low-rank-plus-shift structure.

THEOREM 4.2. *Let $A = [A_1, A_2] \in \mathcal{R}^{m \times n}$ with $m \geq n$. Furthermore assume that*

$$(4.4) \quad A^T A = X + \sigma^2 I, \quad \sigma > 0,$$

where X is symmetric and positive semi-definite with $\text{rank}(X) = k$. Then there are integers $k_1 \leq k$ and $k_2 \leq k$ with $k_1 + k_2 \geq k$ such that

$$(4.5) \quad \text{best}_k([\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]) = \text{best}_k([A_1, A_2]).$$

Proof. The general idea of the proof is to show that what is discarded when A_1 is replaced by $\text{best}_k(A_1)$ and A_2 is replaced by $\text{best}_k(A_2)$ will also be discarded when $\text{best}_k([A_1, A_2])$ is computed from $[A_1, A_2]$. To this end write

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

Since $\text{rank}(X) = k$, it follows that $\text{rank}(A_1^T A_1 - \sigma^2 I) \leq k$ and $\text{rank}(A_2^T A_2 - \sigma^2 I) \leq k$. Let the eigendecompositions of

$$A_1^T A_1 - \sigma^2 I = V_{A_1} \text{diag}(\Sigma_{A_1}^2, 0) V_{A_1}^T, \quad A_2^T A_2 - \sigma^2 I = V_{A_2} \text{diag}(\Sigma_{A_2}^2, 0) V_{A_2}^T,$$

where $\Sigma_{A_1} \in \mathcal{R}^{k_1 \times k_1}$, $\Sigma_{A_2} \in \mathcal{R}^{k_2 \times k_2}$ are nonsingular with $k_1 \leq k$, $k_2 \leq k$. We can write the SVD of A_1 and A_2 as follows:

$$(4.6) \quad A_1 = U_{A_1} \text{diag}(\hat{\Sigma}_{A_1}, \sigma I_{t_1}) V_{A_1}^T = [U_{A_1}^{(1)}, U_{A_1}^{(2)}] \text{diag}(\hat{\Sigma}_{A_1}, \sigma I_{t_1}) [V_{A_1}^{(1)}, V_{A_1}^{(2)}]^T,$$

$$(4.7) \quad A_2 = U_{A_2} \text{diag}(\hat{\Sigma}_{A_2}, \sigma I_{t_2}) V_{A_2}^T = [U_{A_2}^{(1)}, U_{A_2}^{(2)}] \text{diag}(\hat{\Sigma}_{A_2}, \sigma I_{t_2}) [V_{A_2}^{(1)}, V_{A_2}^{(2)}]^T,$$

where $\hat{\Sigma}_{A_1} = (\Sigma_{A_1}^2 + \sigma^2 I_{k_1})^{1/2}$ and $\hat{\Sigma}_{A_2} = (\Sigma_{A_2}^2 + \sigma^2 I_{k_2})^{1/2}$, and $U_{A_1}^{(1)} \in \mathcal{R}^{m \times k_1}$, $U_{A_2}^{(1)} \in \mathcal{R}^{m \times k_2}$, and $t_1 = n_1 - k_1$, $t_2 = n_2 - k_2$, respectively, where n_i is the column dimension of A_i , $i = 1, 2$. Now write $V_{A_1}^T A_1^T A_2 V_{A_2}$ in a partitioned form as

$$(4.8) \quad V_{A_1}^T A_1^T A_2 V_{A_2} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix}, \quad S_{11} \in \mathcal{R}^{k_1 \times k_2}.$$

Since $X = A^T A - \sigma^2 I$ is symmetric positive semi-definite and $\text{rank}(X) = k$, it follows that $S_{12} = 0$, $S_{21} = 0$, $S_{22} = 0$ and $k_1 + k_2 \geq \text{rank}(X) = k$. Using the SVD of A_1 and A_2 in (4.6) and (4.7), Equation (4.8) becomes

$$[U_{A_1}^{(1)} \Sigma_{A_1}, \sigma U_{A_1}^{(2)}]^T [U_{A_2}^{(1)} \Sigma_{A_2}, \sigma U_{A_2}^{(2)}] = \begin{bmatrix} S_{11} & 0 \\ 0 & 0 \end{bmatrix},$$

which leads to²

$$U_{A_1}^{(1)} \perp U_{A_2}^{(2)}, \quad U_{A_1}^{(2)} \perp U_{A_2}^{(1)}, \quad U_{A_1}^{(2)} \perp U_{A_2}^{(2)}.$$

Let \hat{U} be an orthonormal basis of $\mathcal{R}([U_{A_1}^{(1)}, U_{A_2}^{(1)}])$, where $\mathcal{R}(\cdot)$ to denote the column space of a matrix. Then we can write

$$[A_1, A_2] = [\hat{U}, U_{A_1}^{(2)}, U_{A_2}^{(2)}] \text{diag}(\tilde{B}, \sigma I_{t_1}, \sigma I_{t_2}) \begin{bmatrix} (V_{A_1}^{(1)})^T & 0 \\ 0 & (V_{A_2}^{(1)})^T \\ (V_{A_1}^{(2)})^T & 0 \\ 0 & (V_{A_2}^{(2)})^T \end{bmatrix},$$

where $\tilde{B} \in \mathcal{R}^{\hat{k} \times (k_1 + k_2)}$ with all its singular values greater than σ , and $k \leq \hat{k} \leq k_1 + k_2$. Therefore,

$$[A_1, A_2] = \hat{U} \tilde{B} \begin{bmatrix} (V_{A_1}^{(1)})^T & 0 \\ 0 & (V_{A_2}^{(1)})^T \end{bmatrix} + \sigma [U_{A_1}^{(2)}, U_{A_2}^{(2)}] \begin{bmatrix} (V_{A_1}^{(2)})^T & 0 \\ 0 & (V_{A_2}^{(2)})^T \end{bmatrix}$$

the first term in the right hand side of the above is easily seen to be the matrix $[\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)]$, and the relation in Equation (4.5) therefore holds. \square

REMARK. Generically we will have $k_1 + k_2 = k$. In the following we give an example that shows the possibility of the case $k_1 + k_2 > k$. Given any two positive numbers a and b , choose θ_1 and θ_2 such that

$$\cos(\theta_1 - \theta_2) = \sqrt{ab/(a + \sigma^2)(b + \sigma^2)}.$$

Construct two matrices U_1 and U_2 as follows,

$$U_1 = \begin{bmatrix} c_1 & 0 \\ s_1 & 0 \\ 0 & -1 \\ 0 & 0 \end{bmatrix}, \quad U_2 = \begin{bmatrix} c_2 & 0 \\ s_2 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix},$$

where $c_i = \cos(\theta_i)$ and $s_i = \sin(\theta_i)$ for $i = 1, 2$. Now construct the matrix $A = [A_1, A_2]$ with

$$A_1 = U_1 \text{diag}(\sqrt{a + \sigma^2}, \sigma), \quad A_2 = U_2 \text{diag}(\sqrt{b + \sigma^2}, \sigma)$$

we obtain that

$$A^T A - \sigma^2 I = \begin{bmatrix} a & 0 & \sqrt{ab} & 0 \\ 0 & 0 & 0 & 0 \\ \sqrt{ab} & 0 & b & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

and therefore $k = 1$. However, we also have

$$A_1^T A_1 - \sigma^2 I = \begin{bmatrix} a & 0 \\ 0 & 0 \end{bmatrix}, \quad A_2^T A_2 - \sigma^2 I = \begin{bmatrix} b & 0 \\ 0 & 0 \end{bmatrix},$$

² we use $S \perp T$ to denote $S^T T = 0$.

and thus $k_1 = k_2 = 1$. So we have the case $k_1 + k_2 > k$.

REMARK. In essence the result in Theorem 4.2 states that if A has the low-rank-plus-shift structure, then an optimal low-rank approximation of A can be computed by merging the optimal low-rank approximations of its two submatrices A_1 and A_2 . The result can be generalized to the case where A is partitioned into several blocks $A = [A_1, A_2, \dots, A_s]$.

REMARK. In general k_1 and k_2 are not available: they exist in the analysis in the proof of Theorem 4.2 but never explicitly computed. However, since $k_i \leq k, i = 1, 2$, the relation in Equation (4.5) still holds if we replace $k_i, i = 1, 2$, by k , i.e.,

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

Referring back to our discussion on LSI updating problem in Section 2, we see that Theorem 4.2 justifies the replacement of A by its best rank- k approximation because

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

assuming $[A, D]$ has the low-rank-plus-shift structure. That is to say, we will obtain the same best low-rank approximation even though A is replaced by $\text{best}_k(A)$. Numerical results conducted on several text collections show that no retrieval accuracy degradation occurs when updating is computed using a proper implementation [14].

On the other hand, Theorem 4.2 also leads to some novel approaches for computing a low-rank approximation of a large matrix. There are at least two general approaches to pursue ideas based on Theorem 4.2:

- AN INCREMENTAL METHOD. One is what we call incremental approach whereby we can use certain sampling methods to divide the whole collection of documents into several groups: Start with one group and compute its rank- k approximation, and then add the second group using the updating algorithm to produce a new rank- k approximation, and repeat the whole process. This incremental process can be very useful when the data collection is very large and the whole term-document matrix can not reside completely in main memory. Some computational results of this approach can be found in [14].
- A DIVIDE-AND-CONQUER METHOD. Another approach is what we call a divide-and-conquer approach, we can again divide the whole collection of documents into several groups, and compute the rank- k approximation for *each group* and then combine the results together into a rank- k approximation for the whole data collection. Recursively, the rank- k approximation for each group can also be computed using this divide-and-conquer approach and so on. The approach has the property that computation can be organized with high degree of coarse-grain parallelism. A parallel implementation of this method is currently under investigation.

5. Perturbation Analysis. In this section, we consider the case where A only approximately satisfies the low-rank-plus-shift property. Our main goal is to see to what extent the result in Theorem 4.2 still holds in the presence of perturbation. We first present some lemmas which are of their own interests as well. In the sequel $\|\cdot\|$ denotes two-norm and $\|\cdot\|_F$ denotes Frobenius norm. We will use MATLAB notation for submatrices: $A(i:j, k:l)$ denotes rows i to j and columns k to l of A .

LEMMA 5.1. Assume that the matrix X defined below is symmetric positive semi-definite,

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

Then we have

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

Proof. Without loss of generality we assume that the matrix B is diagonal. (The result still holds even if B is rectangular.) Write

$$B = \text{diag}(\sigma_1, \dots, \sigma_n), \quad \sigma_1 \geq \dots \geq \sigma_n.$$

Let a_{11} and c_{11} be the (1,1) element of A and C , respectively. Then

$$\begin{bmatrix} a_{11} & \sigma_1 \\ \sigma_1 & c_{11} \end{bmatrix} = X_1 + E_1,$$

for some 2-by-2 symmetric positive semi-definite X_1 and E_1 with $\|E_1\| \leq \|E\|$. Since the smallest eigenvalue of $X_1 + E_1$ is no smaller than $-\|E_1\|$, it follows that

$$\begin{aligned} \sigma_1^2 &\leq ((a_{11} + c_{11})/2 + \|E_1\|)^2 - ((a_{11} - c_{11})/2)^2 \\ &\leq (\|A\| + \|E\|)(\|C\| + \|E\|), \end{aligned}$$

thus completing the proof. \square

LEMMA 5.2. Let the matrix X be partitioned as

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

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

Proof. The proof is straightforward and is therefore omitted. \square

THEOREM 5.3. Let $A = [A_1, A_2] \in \mathcal{R}^{m \times n}$, $m \geq n$ satisfy

$$A^T A = X + \sigma^2 I + E,$$

where X symmetric positive semi-definite with $\text{rank}(X) = k$. If

$$\lambda_k(X) > 3\|E\| + 2\sqrt{\|E\|(\|X\| + 2\|E\|)},$$

then for some $k_1 \leq k$, $k_2 \leq k$, and $k_1 + k_2 \geq k$, we have

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

with

$$\|\Delta\| \leq \alpha(\|E\| + 4\sqrt{\|E\|(\|X\| + 2\|E\|)})$$

and

$$\alpha = \sqrt{\lambda_k(X) - \|E\|}/\lambda_k(X).$$

Proof. The proof is divided into several parts.

1) We first write the eigendecomposition of the following matrices:

$$A^T A - \sigma^2 I = X + E = V \text{diag}(\lambda_1, \dots, \lambda_n) V^T,$$

$$A_i^T A_i - \sigma^2 I = X_i + E_i = V_i \text{diag}(\lambda_1^{(i)}, \dots, \lambda_{n_i}^{(i)}) V_i^T, \quad i = 1, 2,$$

where X and X_i are symmetric positive semi-definite, and V and V_i are orthogonal for $i = 1, 2$. The eigenvalues $\{\lambda_k\}$ and $\{\lambda_k^{(i)}\}$ are arranged in nonincreasing order. It follows that, for $i = 1, 2$, there are orthogonal matrices U_i such that

$$\begin{aligned} A_i &= U_i \text{diag}(d_1^{(i)}, \dots, d_{n_i}^{(i)}) V_i^T, \\ &= U_i \text{diag}(D_{i1}, D_{i2}) V_i^T, \end{aligned}$$

where $d_j^{(i)} = \sqrt{\lambda_j^{(i)} + \sigma^2}$ and

$$\begin{aligned} D_{i1} &= \text{diag}(d_1^{(i)}, \dots, d_{k_i}^{(i)}), \\ D_{i2} &= \text{diag}(d_{k_i+1}^{(i)}, \dots, d_{n_i}^{(i)}), \end{aligned}$$

where $k_i = \text{rank}(X_i)$, $i = 1, 2$. The definition of best low-rank approximation leads to

$$\text{best}_{k_i}(A_i) = U_i D_{i1} V_i, \quad i = 1, 2.$$

2) Using the above decompositions we now write the matrix A in several different forms: $A = [A_1, A_2] = B[W_1, W_2]^T = [B_1, B_2]W^T$, where

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

and

$$W_1 = \begin{bmatrix} V_{11} & 0 \\ 0 & V_{21} \end{bmatrix}, \quad W_2 = \begin{bmatrix} V_{12} & 0 \\ 0 & V_{22} \end{bmatrix}.$$

It can be readily verified that

$$B_1 W_1^T = [\text{best}_{k_1}(A_1), \text{best}_{k_2}(A_2)].$$

Now partition

$$(5.9) \quad B^T B - \sigma^2 I = \begin{bmatrix} B_1^T B_1 - \sigma^2 I & B_1^T B_2 \\ B_2^T B_1 & B_2^T B_2 - \sigma^2 I \end{bmatrix}.$$

3) Let the eigendecomposition of $B_1^T B_1 - \sigma^2 I$ be

$$B_1^T B_1 - \sigma^2 I = G \text{diag}(\alpha_1, \dots, \alpha_s) G^T,$$

where $\alpha_1 \geq \dots \geq \alpha_s$, $s = k_1 + k_2$. Now partition $\text{diag}(G^T, I) B^T B \text{diag}(G, I)$ as

$$\text{diag}(G^T, I) B^T B \text{diag}(G, I) = \begin{bmatrix} C & E_1^T \\ E_1 & E_2 \end{bmatrix},$$

where $C = \text{diag}(\alpha_1, \dots, \alpha_k)$, and the matrix E has the form

$$(5.10) \quad [E_1, E_2] = \begin{bmatrix} [0, \text{diag}(\alpha_{k+1}, \dots, \alpha_s)] & (G^T B_1^T B_2)(k+1:m, :) \\ B_2^T B_1 G & B_2^T B_2 - \sigma^2 I \end{bmatrix}.$$

Furthermore, let the eigendecomposition of

$$\begin{bmatrix} C & E_1^T \\ E_1 & E_2 \end{bmatrix} = Q \Lambda Q^T, \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_n).$$

Now partition conformally,

$$Q = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix}, \quad \Lambda = \text{diag}(\Lambda_1, \Lambda_2)$$

with $\Lambda_1 = \text{diag}(\lambda_1, \dots, \lambda_k)$. It can be verified that

$$(5.11) \quad Q_{21} = [E_1, E_2] \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix} \Lambda_1^{-1}.$$

It follows that we can find U orthogonal such that

$$B = [B_1, B_2] = U(\Lambda + \sigma^2 I)^{1/2} Q^T \text{diag}(G^T, I),$$

which leads to

$$B_1 W_1^T = U(\Lambda + \sigma^2 I)^{1/2} Q^T \begin{bmatrix} G^T W_1^T \\ 0 \end{bmatrix},$$

and hence

$$\text{best}_k(B_1 W_1^T) = U(:, 1:k)(\Lambda_1 + \sigma^2 I)^{1/2} (Q(:, 1:k))^T \begin{bmatrix} G^T W_1^T \\ 0 \end{bmatrix}.$$

4) On the other hand, from

$$A = U(\Lambda + \sigma^2 I)^{1/2} Q^T \begin{bmatrix} G^T W_1^T \\ W_2^T \end{bmatrix}$$

it follows that

$$\text{best}_k(A) = U(:, 1:k)(\Lambda_1 + \sigma^2 I)^{1/2} (Q(:, 1:k))^T \begin{bmatrix} G^T W_1^T \\ W_2^T \end{bmatrix} = \text{best}_k(B_1 W_1^T) + \Delta,$$

where

$$\Delta = U(:, 1:k)(\Lambda_1 + \sigma^2 I)^{1/2} Q_{21}^T W_2^T$$

with

$$\|\Delta\| \leq \|(\Lambda_1 + \sigma^2 I)^{1/2} Q_{21}^T\| \leq \|[E_1, E_2]\| \sqrt{\lambda_k + \sigma^2 / \lambda_k},$$

where we have used Equation (5.11). Now we need to bound $\|[E_1, E_2]\|$, and this will be done in the following steps.

5) Applying Lemma 5.1 to the matrix

$$B_2^T B_2 - \sigma^2 I = \begin{bmatrix} D_{12}^2 - \sigma^2 I & D_{11}^T U_{12}^T U_{22} D_{22} \\ D_{22}^T U_{22}^T U_{12} D_{11} & D_{22}^2 - \sigma^2 I \end{bmatrix},$$

we obtain the bound

$$\|D_{11}^T U_{12}^T U_{22} D_{22}\|^2 \leq (\|D_{12}^2 - \sigma^2 I\| + \|E\|)(\|D_{22}^2 - \sigma^2 I\| + \|E\|) \leq 4\|E\|^2,$$

where we have used $\|D_{i2}^2 - \sigma^2 I\| \leq \|E\|, i = 1, 2$. By Lemma 5.2 we obtain

$$\|B_2^T B_2 - \sigma^2 I\| \leq 3\|E\|.$$

6) Now applying Lemma 5.1 to $B^T B - \sigma^2 I$ in Equation (5.9) we obtain

$$\|B_2^T B_1\|^2 \leq 4\|E\|(\|B_1^T B_1 - \sigma^2 I\| + \|E\|) \leq 4\|E\|(\|X\| + 2\|E\|),$$

where we have bounded

$$\|B_1^T B_1 - \sigma^2 I\| \leq \|B^T B - \sigma^2 I\| \leq \|X\| + \|E\|.$$

7) Using Equation (5.10) and the results in Part 5) and 6) we have

$$\|[E_1, E_2]\| \leq \max\{|\alpha_{k+1}|, \dots, |\alpha_{k+1}|, 3\|E\|\} + 2\sqrt{\|E\|(\|X\| + 2\|E\|)}.$$

8) Now we are ready to complete the proof of the theorem by showing that for $j > k$,

$$|\alpha_j| \leq \|E\| + 2\sqrt{\|E\|(\|X\| + 2\|E\|)}.$$

In fact, by definition, $\lambda_j = \lambda_j(B^T B - \sigma^2 I)$, and therefore

$$|\lambda_j - \lambda_j(\text{diag}(B_1^T B_1 - \sigma^2 I, B_2^T B_2 - \sigma^2 I))| \leq \|B_2^T B_1\|.$$

On the other hand, we have for $j > k$,

$$|\lambda_j(B_2^T B_2 - \sigma^2 I)| \leq \|B_2^T B_2 - \sigma^2 I\| \leq 3\|E\|.$$

The assumption of the theorem implies that

$$\lambda_k > \|B_2^T B_1\| + 3\|E\|,$$

and therefore for $j \leq k$,

$$\lambda_j(\text{diag}(B_1^T B_1 - \sigma^2 I, B_2^T B_2 - \sigma^2 I)) = \lambda_j(B_1^T B_1 - \sigma^2 I) = \alpha_j.$$

Now for any $j > k$, there is $i_j > k$ such that $|\alpha_j - \lambda_{i_j}| \leq \|B_2^T B_1\|$, and thus

$$|\alpha_j| \leq |\lambda_{i_j}| + \|B_2^T B_1\| \leq \|E\| + 2\sqrt{\|E\|(\|X\| + 2\|E\|)},$$

completing the proof. \square

REMARK. In many of our numerical experiments, we observed that $\|\Delta\| = O(\|E\|)$ versus the bound $\|\Delta\| = O(\|E\|^{1/2})$ given in the above theorem. Here we

give an example for which we do have $\|\Delta\| = O(\|E\|^{1/2})$. Let ϵ be small, and define $t = \sqrt{2\epsilon}$. Denote

$$\begin{aligned}\mu_1 &= (1 + \epsilon + \sqrt{(1 + \epsilon)^2 + 4\epsilon})/2 = 1 + 2\epsilon - 2\epsilon^2 + O(\epsilon^3), \\ \mu_2 &= (1 + \epsilon - \sqrt{(1 + \epsilon)^2 + 4\epsilon})/2 = -\epsilon + 2\epsilon^2 + O(\epsilon^3), \\ a &= (\mu_1 - 1)/t = t(1 - \epsilon + O(\epsilon^2)).\end{aligned}$$

Then it can be readily verified that

$$\begin{bmatrix} 1 & t \\ t & \epsilon \end{bmatrix} = (1/(1 + a^2)) \begin{bmatrix} 1 & a \\ a & -1 \end{bmatrix} \begin{bmatrix} \mu_1 & 0 \\ 0 & \mu_2 \end{bmatrix} \begin{bmatrix} 1 & a \\ a & -1 \end{bmatrix}.$$

Now let

$$X = (\mu_1/(1 + a^2)) \begin{bmatrix} 1 & 0 \\ 0 & a \\ 0 & 1 \\ a & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & a \\ 0 & a & 1 & 0 \end{bmatrix},$$

$$E = (\mu_2/(1 + a^2)) \begin{bmatrix} 0 & a \\ 1 & 0 \\ -a & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 0 & 1 & -a & 0 \\ a & 0 & 0 & -1 \end{bmatrix}.$$

With $k = \text{rank}(X) = 2$ and $\|E\| = |\mu_2| = \epsilon(1 - 2\epsilon + O(\epsilon^2))$, it is easy to see that

$$X + E = V \text{diag}(\mu_1, \mu_1, \mu_2, \mu_2)V,$$

where

$$V = (1/(1 + a^2)) \begin{bmatrix} 1 & 0 & 0 & a \\ 0 & a & 1 & 0 \\ 0 & 1 & -a & 0 \\ a & 0 & 0 & -1 \end{bmatrix}$$

is orthogonal and symmetric. For a given $\sigma > 0$, define

$$c_1 = \sqrt{\mu_1 + \sigma^2}, \quad c_2 = \sqrt{\mu_2 + \sigma^2}.$$

Now construct matrix A as $A = [A_1, A_2]$, where

$$A_1 = [e_1, e_3, e_2, e_4]DV(:, 1:2), \quad A_2 = [e_1, e_3, e_2, e_4]DV(:, 3:4)$$

with $D = \text{diag}(c_1, c_1, c_2, c_2)$, and $I_4 = [e_1, e_2, e_3, e_4]$ is the identity matrix. Then we have

$$\text{best}_k(A) = (c_1/\sqrt{1 + a^2})[e_1, e_3]V(1:2, :).$$

Since $a < 1$, it can be verified that

$$\text{best}_1(A_1) = (1/\sqrt{1 + a^2}) \begin{bmatrix} c_1 & 0 \\ 0 & 0 \\ 0 & 0 \\ c_2 a & 0 \end{bmatrix}, \quad \text{best}_1(A_2) = (1/\sqrt{1 + a^2}) \begin{bmatrix} 0 & 0 \\ -c_2 a & 0 \\ c_1 & 0 \\ 0 & 0 \end{bmatrix}.$$

It follows that

$$\text{best}_k([\text{best}_1(A_1), \text{best}_1(A_2)]) = (1/\sqrt{1+a^2}) \begin{bmatrix} c_1 & 0 & 0 & 0 \\ 0 & 0 & -c_2 a & 0 \\ 0 & 0 & c_1 & 0 \\ c_2 a & 0 & 0 & 0 \end{bmatrix},$$

and therefore

$$\begin{aligned} \Delta &= \text{best}_k(A) - \text{best}_k([\text{best}_1(A_1), \text{best}_1(A_2)]) \\ &= a/\sqrt{1+a^2} \text{diag}(c_1, c_2, c_1, c_2)[-e_4, e_3, e_2, e_1] \end{aligned}$$

which leads to $\|\Delta\| = ac_1/\sqrt{1+a^2}$. Then we see that

$$\lim_{\epsilon \rightarrow 0} \|\Delta\|/\sqrt{\|E\|} = \sqrt{2(1+\sigma^2+2\epsilon+O(\epsilon^2))}/[(1+O(\epsilon^2))(1-\epsilon+O(\epsilon^2))] = \sqrt{2(1+\sigma^2)}.$$

6. Concluding Remarks. In this paper we present a detailed analysis of matrices with low-rank-plus-shift structure. Our emphasis is placed on justifying some novel methods for partial SVD computation and partial SVD updating problems arising from LSI in information retrieval. Our perturbation analysis demonstrates that the results we have derived are still valid even the low-rank-plus-shift structure is approximately satisfied. The results we have proved provide theoretical justifications for the novel LSI updating algorithms and the incremental and divide-and-conquer approaches proposed in [14, 16]. Our future research will concentrate on further developing the numerical algorithms and their parallel implementations. We will also refine our perturbation analysis, especially we will try to find conditions on the matrix A that will allow us to improve the perturbation bounds in Theorem 5.3 from $O(\|E\|^{1/2})$ to $O(\|E\|)$.

REFERENCES

- [1] M. Berry. Large Scale Singular Value Computations. *International Journal of Supercomputer Applications*, 6:13-49, 1992.
- [2] M.W. Berry, S.T. Dumais and G.W. O'Brien. Using linear algebra for intelligent information retrieval. *SIAM Review*, 37:573-595, 1995.
- [3] Cornell SMART System, <ftp://ftp.cs.cornell.edu/pub/smart>.
- [4] CMU Data, <http://www.cs.cmu.edu/afs/cs/project/theo-11/www/naive-bayes.html>.
- [5] S. Deerwester, S.T. Dumais, T.K. Landauer, G.W. Furnas and R.A. Harshman. Indexing by latent semantic analysis. *Journal of the Society for Information Science*, 41:391-407, 1990.
- [6] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, Maryland, 2nd edition, 1989.
- [7] G. Kowalski. Information Retrieval System: Theory and Implementation. Kluwer Academic Publishers, Boston, 1997.
- [8] R. Krovetz and W.B. Croft. Lexical ambiguity and information retrieval. *ACM Transactions on Information Systems*, 10:115-141, 1992.
- [9] B.N. Parlett. *The Symmetric Eigenvalue Problem*. Prentice Hall Inc., Englewood Cliffs, New Jersey, 1980.
- [10] G. Salton. Automatic Text Processing. Addison-Wesley, New York, 1989.
- [11] 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.)
- [12] G. Xu and T. Kailath. Fast subspace decomposition. *IEEE Transactions on Signal Processing*, 42:539-551, 1994.
- [13] 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.

- [14] 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*.
- [15] H. Zha and H. Simon. A Subspace-based model for Latent Semantic Indexing in information retrieval. To appear in *Proceedings of Interface '98*.
- [16] H. Zha, O. Marques and H. Simon. Large-scale SVD and Subspace-based Methods for Information Retrieval. To appear in *Proceedings of Irregular '98*, Lecture Notes in Computer Science, Springer-Verlag.
- [17] H. Zha and Z. Zhang. A sorted partial Jacobi method and its convergence analysis. *Linear Algebra and its Applications*, 270:79-108, 1997.

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