

AN ALTERNATIVE ALGORITHM FOR THE REFINEMENT OF ULV DECOMPOSITIONS*

JESSE L. BARLOW[†], HASAN ERBAY[‡], AND IVAN SLAPNIČAR[§]

Abstract. The ULV decomposition (ULVD) is an important member of a class of rank-revealing two-sided orthogonal decompositions used to approximate the singular value decomposition (SVD). It is useful in applications of the SVD such as principal components where we are interested in approximating a matrix by one of lower rank. It can be updated and downdated much more quickly than an SVD. In many instances, the ULVD must be refined to improve the approximation it gives for the important right singular subspaces or to improve the matrix approximation. Present algorithms to perform this refinement require $O(mn)$ operations if the rank of the matrix is k , where k is very close to 0 or n , but these algorithms require $O(mn^2)$ operations otherwise. Presented here is an alternative refinement algorithm that requires $O(mn)$ operations no matter what the rank is. Our tests show that this new refinement algorithm produces similar improvement in matrix approximation and in the subspaces. We also propose slight improvements on the error bounds on subspaces and singular values computed by the ULVD.

Key words. subspace approximation, orthogonal decompositions, refinement

AMS subject classifications. 65F25, 65F15

DOI. 10.1137/S0895479801372621

1. Introduction. The singular value decomposition (SVD) of a matrix $X \in \mathbf{R}^{m \times n}$, $m \geq n$, is written

$$(1.1) \quad X = Y\Sigma W^T,$$

where $W \in \mathbf{R}^{n \times n}$ is orthogonal and $Y \in \mathbf{R}^{m \times n}$ is *left orthogonal*, that is, $Y^T Y = I_n$, and

$$(1.2) \quad \Sigma = \text{diag}(\sigma_1, \dots, \sigma_n) \in \mathbf{R}^{n \times n}, \quad \sigma_1 \geq \dots \geq \sigma_n.$$

It contains useful information about a linear operator, including rank, important subspaces, and conditioning. It also gives us the ability to compute low rank approximations.

Some of the most important applications of the SVD include solving ill-conditioned least squares problems [5], total least squares problems [22], subspace tracking, and isolating principal components [15].

In these applications, our interest in the SVD is to write X in the form

$$(1.3) \quad X = Y_1 \Sigma_1 W_1^T + Y_2 \Sigma_2 W_2^T,$$

*Received by the editors January 31, 2001; accepted for publication (in revised form) by L. Eldén September 28, 2004; published electronically July 13, 2005. This research was supported by the National Science Foundation under contract CCR-9732081.

<http://www.siam.org/journals/simax/27-1/37262.html>

[†]Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA 16802-6106 (barlow@cse.psu.edu).

[‡]Corresponding author. Mathematics Department, Arts and Science Faculty, Kirikkale University, Kirikkale 71100, Turkey (hxe68@yahoo.com).

[§]Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture, University of Split, R. Boškovića b.b., 21000 Split, Croatia (ivan.slapnicar@fesb.hr).

where

$$\Sigma_1 = \text{diag}(\sigma_1, \dots, \sigma_k), \quad \Sigma_2 = \text{diag}(\sigma_{k+1}, \dots, \sigma_n)$$

and $Y_1, Y_2, W_1,$ and W_2 are the corresponding matrices of left and right singular vectors. The other information from the SVD is not needed. Moreover, accurate representations of $\text{Range}(W_1)$ and/or $\text{Range}(W_2)$ are often satisfactory. In principal component analysis, we only need a good low rank approximation of X .

We use $\|\cdot\|$ to denote the two norm (matrix and vector), and $\|\cdot\|_F$ to denote the Frobenius norm. The notation $\sigma_k(X)$ denotes the k th largest singular value of X .

Complete orthogonal decompositions (CODs) can compute an approximation to the dichotomy in (1.3) slightly faster than an SVD. However, if the decomposition must be modified by adding rows (updating) and/or deleting rows (downdating), then the SVD can be computationally expensive [13], whereas CODs may be updated and downdated in $O(mn)$ flops [8, 3, 24, 20, 21]. Our algorithms for construction and refinement lead to a form that can be updated by any of these procedures.

The CODs have the form

$$(1.4) \quad X = UCV^T,$$

where C is a triangular matrix partitioned according to

$$C = \begin{pmatrix} k & n-k \\ C_1 & C_2 \end{pmatrix}, \quad \|C_2\| \approx \|\Sigma_2\| = \sigma_{k+1}$$

and $U \in \mathbf{R}^{m \times n}$ is left orthogonal, $V \in \mathbf{R}^{n \times n}$ is orthogonal.

The first such decompositions were proposed independently in [9] and in [14]. These used maximal column pivoting [6]. More recent enhancements, such as the URV and ULV decompositions, discussed in [20, 21], made use of condition estimators. This paper uses the ULV decomposition (ULVD), which we define in the form (1.4), where, for a fixed integer k and tolerance ϵ , we have

$$(1.5) \quad U = \begin{pmatrix} k & m-k \\ U_1 & U_2 \end{pmatrix}, \quad V = \begin{pmatrix} k & n-k \\ V_1 & V_2 \end{pmatrix},$$

$$(1.6) \quad C = \begin{pmatrix} k & n-k \\ n-k & \begin{pmatrix} L & 0 \\ F & G \end{pmatrix} \end{pmatrix},$$

and

$$(1.7) \quad \|L^{-1}\|^{-1} \leq \epsilon, \quad \|(F \ G)\| \leq \Phi(n)\epsilon.$$

Here $\Phi(n)$ is a modestly growing function of n , say, \sqrt{n} . This formulation is a variant of the one in [24, 8, 4]. We refer to k as the ϵ -pseudorank of X , corresponding to a definition used by Lawson and Hanson [16, p. 77].

The matrix V yields approximations to the two right singular subspaces, $\text{Range}(W_1)$ and $\text{Range}(W_2)$, the matrix U yields approximations to the two left singular subspaces, $\text{Range}(Y_1)$ and $\text{Range}(Y_2)$. The accuracy of these approximations and approximations to the singular values of C are discussed in section 2. The results given there are slight enhancements of those by Mathias and Stewart [17] and Fierro and Bunch [10].

In this work, we give improvements to a refinement procedure for the ULVD in [17]. Often after an update or downdate, the conditions (1.7) are not satisfied,

and/or $\|(F \ G)\|_F$ is not small enough. Thus procedures which make $\|(F \ G)\|_F$ smaller are necessary. Procedures are given in [17, 10], but these require $O(mn^2)$ flops (floating point operations; see [12, pp. 18–19]) unless k is very close to 0 or n . Our new procedure (Algorithm 4.2) requires $O(mn)$ flops regardless of the value of k .

Throughout the paper, we use MATLAB notation for submatrices; thus $X(i: j, :)$ denotes rows i through j of X , $X(:, k: \ell)$ denotes columns k through ℓ , and $X(i: j, k: \ell)$ denotes the intersection of the two.

In the next section, we give some improvements of singular value and vector bounds in [17, 10]. In section 3, we introduce the necessary matrix computational tools to build the algorithms for the ULVD. Our new refinement procedure (Algorithm 4.2) is given in section 4 along with some results on its properties. Numerical tests of the algorithm given in section 5 show that the new algorithm is much faster than the one in [17], but just as effective. A conclusion is given in section 6.

2. Subspaces and singular values. The ULVD (1.4)–(1.5) generates approximations to the right and left singular subspaces associated with the first k and last $n - k$ singular values. In the Davis–Kahan [7] framework the errors in these subspaces are characterized by

$$(2.1) \quad |\sin \Theta_R| = \|W_1^T V_2\|_2 = \|V_1^T W_2\|_2,$$

$$(2.2) \quad |\sin \Theta_L| = \|Y_1^T U_2\|_2 = \|U_1^T Y_2\|_2.$$

The following proposition is a slight improvement on bounds by Mathias and Stewart [17] and Fierro and Bunch [10]. It characterizes the error in the singular subspaces and singular values. The long but elementary proof is in the longer technical report version of this paper [2].

PROPOSITION 2.1. *Let X have the ULVD (1.4)–(1.5) and let*

$$(2.3) \quad \rho \stackrel{\text{def}}{=} \min \left\{ \frac{\sigma_{k+1}(X)}{\sigma_k(L)}, \frac{\sigma_1(G)}{\sigma_k(X)} \right\},$$

$$(2.4) \quad \gamma_R \stackrel{\text{def}}{=} \max \left\{ \sigma_k^2(X) - \sigma_1^2(G), \sigma_k^2 \left[\begin{pmatrix} L \\ F \end{pmatrix} \right] - \sigma_{k+1}^2(X) \right\},$$

$$(2.5) \quad \gamma_L \stackrel{\text{def}}{=} \max \left\{ \frac{\sigma_k^2(X) - \sigma_1^2(G)}{\sigma_k(X)}, \frac{\sigma_k^2(L) - \sigma_{k+1}^2(X)}{\sigma_k(L)} \right\}.$$

Then

$$(2.6) \quad |\sin \Theta_R| \leq \frac{\|F^T G\|_2}{\gamma_R}, \quad |\sin \Theta_L| \leq \frac{\|F\|_2}{\gamma_L},$$

$$(2.7) \quad |\sin \Theta_R| \leq \rho |\sin \Theta_L|,$$

$$(2.8) \quad |\cos \Theta_L| \leq \frac{\sigma_j(L)}{\sigma_j(X)}, \frac{\sigma_{i+k}(X)}{\sigma_i(G)} \leq 1, \quad \begin{array}{l} j = 1, 2, \dots, k, \\ i = 1, 2, \dots, n - k. \end{array}$$

The bounds in [17] and [10] are the right side of the following two inequalities; thus the bounds in (2.6) are sharper.

$$|\sin \Theta_R| \leq \frac{\|F^T G\|_2}{\gamma_R} \leq \frac{\|F\|_2 \|G\|_2}{\sigma_k^2(L) - \sigma_1^2(G)},$$

$$|\sin \Theta_L| \leq \frac{\|F\|_2}{\gamma_L} \leq \frac{\|F\|_2 \sigma_k(L)}{\sigma_k^2(L) - \sigma_1^2(G)}.$$

Through the use of trigonometric identities, (2.8) implies the relative error bound

$$(2.9) \quad 0 \leq \frac{\sigma_j(X) - \sigma_j(L)}{\sigma_j(X)}, \frac{\sigma_i(G) - \sigma_{i+k}(X)}{\sigma_{i+k}(X)} \leq \frac{\sin^2 \Theta_L}{1 + |\cos \Theta_L|}$$

for $j = 1, 2, \dots, k$ and $i = 1, 2, \dots, n - k$. Thus as pointed out in [17], the relative error in the singular values is proportional to $\|F\|_2^2$.

3. Matrix computational tools for the ULVD.

3.1. Norm and condition estimates. For a lower triangular matrix $C \in \mathbf{R}^{n \times n}$, we need to estimate its smallest singular value and associated vectors given by the triple $(\sigma_n, \mathbf{u}_n, \mathbf{v}_n)$. Realistically, we need to compute $\bar{\sigma}_n$, $\bar{\mathbf{u}}_n$, and $\bar{\mathbf{v}}_n$ such that

$$(3.1) \quad C^{-1} \bar{\mathbf{u}}_n = \bar{\sigma}_n^{-1} \bar{\mathbf{v}}_n,$$

$$(3.2) \quad C^{-T} \bar{\mathbf{v}}_n = \bar{\sigma}_n^{-1} \bar{\mathbf{u}}_n + \mathbf{r}_{inv},$$

$$(3.3) \quad \mathbf{r}_{inv}^T \bar{\mathbf{u}}_n = 0, \quad \|\mathbf{r}_{inv}\| \leq \bar{\sigma}_n^{-1} \delta,$$

where δ is some prescribed tolerance which should be smaller than ϵ in (1.5). In our tests in section 5, we used $\delta = 10^{-3}\epsilon$. In practice, we expect to use only a few Lanczos or power iterations and some back and forward solves with C to perform this computation.

Likewise, to estimate the largest singular value and associated vectors, given by the triple $(\sigma_1, \mathbf{u}_1, \mathbf{v}_1)$, we compute a triple $(\bar{\sigma}_1, \bar{\mathbf{u}}_1, \bar{\mathbf{v}}_1)$ such that

$$(3.4) \quad C \bar{\mathbf{v}}_1 = \bar{\sigma}_1 \bar{\mathbf{u}}_1,$$

$$(3.5) \quad C^T \bar{\mathbf{u}}_1 = \bar{\sigma}_1 \bar{\mathbf{v}}_1 + \mathbf{r},$$

$$(3.6) \quad \mathbf{r}^T \bar{\mathbf{v}}_1 = 0, \quad \|\mathbf{r}\| \leq \bar{\sigma}_1 \delta.$$

Again, this can be achieved with a few Lanczos or power iterations. Here we do not restrict C to be lower triangular or even square. For complexity estimates, we assume a fixed number of iterations, c_{iter} , for both of these procedures.

When these operations are counted in complexity estimates, we count (3.1)–(3.3) as $c_{iter}n^2$ flops, and for $C \in \mathbf{R}^{m \times n}$ we count (3.4)–(3.6) as $2c_{iter}mn$ flops.

When using these procedures in the analysis of other algorithms, we ignore the errors \mathbf{r} and \mathbf{r}_{inv} , though they affect the accuracy.

3.2. Chasing algorithms. Once we have computed the approximate singular vectors in section 3.1, we need the two chasing procedures described in this section. Such algorithms are discussed in [21, 10, 1, 4], so we just summarize them here.

First, suppose that we have computed $(\bar{\sigma}_n, \bar{\mathbf{u}}_n, \bar{\mathbf{v}}_n)$ from (3.1)–(3.3). To “deflate” this smallest singular value, we find an orthogonal matrix \bar{U} such that

$$(3.7) \quad \bar{U}^T \bar{\mathbf{u}}_n = \pm \mathbf{e}_n$$

and an orthogonal matrix \bar{V} such that

$$(3.8) \quad \bar{U}^T C \bar{V} = \hat{C}$$

is lower triangular. The matrices \bar{U} and \bar{V} are each the products of $n - 1$ Givens rotations.

The algorithm (3.7)–(3.8) has interesting consequences for \hat{C} . If

$$(3.9) \quad \hat{\mathbf{v}}_n = \bar{V}^T \bar{\mathbf{v}}_n, \quad \hat{\mathbf{r}}_{inv} = U^T \mathbf{r}_{inv},$$

then (3.1)–(3.3) becomes

$$(3.10) \quad \hat{C} \hat{\mathbf{v}}_n = \pm \bar{\sigma}_n \mathbf{e}_n,$$

$$(3.11) \quad \pm \hat{C}^T \mathbf{e}_n = \bar{\sigma}_n \hat{\mathbf{v}}_n + \bar{\sigma}_n \hat{C}^T \hat{\mathbf{r}}_{inv}, \quad \|\hat{\mathbf{r}}_{inv}\| \leq \delta \bar{\sigma}_n^{-1}.$$

It is easily concluded from (3.10) that

$$\begin{aligned} \hat{\mathbf{v}}_n &= \pm \mathbf{e}_n, & \hat{c}_{nn} &= \bar{\sigma}_n, \\ \|\hat{C}^T \mathbf{e}_n\|^2 &= \bar{\sigma}_n^2 + \bar{\sigma}_n^2 \|\hat{C}^T \hat{\mathbf{r}}_{inv}\|^2 \leq \bar{\sigma}_n^2 + \|\hat{C}\|^2 \delta^2 \end{aligned}$$

and that

$$(0 \ \hat{C}(n, 2 : n)) = \mp \bar{\sigma}_n \hat{\mathbf{r}}_{inv}^T \hat{C}.$$

Thus

$$\|\hat{C}(n, 2 : n)\| \leq \delta \|C\|.$$

Similarly, if C is lower trapezoidal with more columns than rows, we can develop a chasing procedure similar to that above that starts with $\bar{\mathbf{u}}_1$ in (3.4)–(3.6) and produces orthogonal matrices \bar{U} and \bar{V} such that

$$(3.12) \quad \bar{U}^T \bar{\mathbf{u}}_1 = \pm \mathbf{e}_1,$$

and \bar{V} such that \hat{C} in (3.8) remains lower trapezoidal.

If $\hat{\mathbf{v}}_1 = V^T \bar{\mathbf{v}}_1$ and $\hat{\mathbf{r}} = V^T \mathbf{r}$, then

$$\begin{aligned} \hat{C} \hat{\mathbf{v}}_1 &= \pm \bar{\sigma}_1 \mathbf{e}_1, \\ \pm \hat{C}^T \mathbf{e}_1 &= \bar{\sigma}_1 \hat{\mathbf{v}}_1 + \hat{\mathbf{r}}, \quad \|\hat{\mathbf{r}}\| \leq \delta \bar{\sigma}_1. \end{aligned}$$

Since $\hat{\mathbf{r}}^T \hat{\mathbf{v}}_1 = 0$, we have

$$(3.13) \quad \|\hat{C}^T \mathbf{e}_1\|^2 = \bar{\sigma}_1^2 + \|\hat{\mathbf{r}}\|^2,$$

and thus

$$(3.14) \quad \bar{\sigma}_1 \leq \|\hat{C}^T \mathbf{e}_1\| \leq \bar{\sigma}_1 (1 + \delta^2)^{1/2}.$$

4. Refining a ULVD. Refinement is a process of improving the accuracy of the ULVD as an approximation to the SVD. We discuss how to measure that improvement below.

The notion of refinement was popularized in [17], where the following was presented. It is equivalent to the QR algorithm without shifts.

ALGORITHM 4.1 (Mathias–Stewart refinement procedure).

Step 1. Find orthogonal $\tilde{U} \in \mathbf{R}^{n \times n}$ such that

$$\begin{aligned} \tilde{U}^T C &= \tilde{U}^T \begin{pmatrix} L & 0 \\ F & G \end{pmatrix} = \begin{pmatrix} L^{(1)} & F^{(1)} \\ 0 & G^{(1)} \end{pmatrix} = C^{(1)}, \\ U &\leftarrow U \tilde{U}. \end{aligned}$$

Step 2. Find orthogonal $\tilde{V} \in \mathbf{R}^{n \times n}$ such that

$$C^{(1)}\tilde{V} = \begin{pmatrix} L^{(1)} & F^{(1)} \\ 0 & G^{(1)} \end{pmatrix} \tilde{V} = \begin{pmatrix} \bar{L} & 0 \\ \bar{F} & \bar{G} \end{pmatrix},$$

$$V \leftarrow V\tilde{V}.$$

We note three important properties of Algorithm 4.1. Other properties were shown in [17].

(1) Since this is the QR algorithm without shifts, it follows that asymptotically

$$\|\bar{F}\| \approx \|F\| \frac{\sigma_{k+1}^2(C)}{\sigma_k^2(C)}.$$

(2) One can show that

$$\|(\bar{F} \ \bar{G})\| \leq \|G\|, \quad \|(\bar{F} \ \bar{G})\|_F \leq \|G\|_F.$$

(3) Repeated application of this algorithm forces $\|(\bar{F} \ \bar{G})\|$ to converge to $\sigma_{k+1}(C)$, but at about the rate of the shiftless QR algorithm.

This algorithm requires $O((m+n)(n-k)k)$ flops, where the constant in $O(\cdot)$ depends upon the implementation. Thus, if $\min\{k, n-k\} = O(1)$, then this is an $O(mn)$ algorithm, but if $\min\{k, n-k\} = O(n)$, this is an $O(mn^2)$ algorithm. In our tests, we used a chasing procedure that combines Steps 1 and 2.

An alternative to Algorithm 4.1, proposed in [10], just computes a ULVD of C , and thereby produce a better ULVD of X . This costs about the same as Algorithm 4.1— $O(mn(n-k))$ flops—making it $O(mn)$ flops if $n-k = O(1)$ and $O(mn^2)$ flops if $n-k = O(n)$.

Below we give an algorithm that computes a refinement in $O(mn)$ flops no matter what the value of k is. For our statement of this refinement procedure, C has the form (1.6).

ALGORITHM 4.2 (alternative refinement algorithm).

Step 1. Using procedures such as those in section 3.1, find unit vectors $\bar{\mathbf{u}}_1, \bar{\mathbf{v}}_1$ such that

$$(F \ G) \bar{\mathbf{v}}_1 = \bar{\sigma}_1 \bar{\mathbf{u}}_1,$$

where $\bar{\sigma}_1$ is the largest singular value of $(F \ G)$.

Step 2. Using the method discussed in section 3.2, construct orthogonal matrices $U^{(1)}$ and $V^{(1)}$ such that

$$[U^{(1)}]^T \bar{\mathbf{u}}_1 = \pm \mathbf{e}_1,$$

and

$$G^{(1)} \leftarrow [U^{(1)}]^T G V^{(1)}$$

remains lower triangular. Also compute

$$F^{(1)} \leftarrow [U^{(1)}]^T F.$$

Note that

$$(F^{(1)} \ G^{(1)}) = \frac{1}{n-k-1} \begin{pmatrix} k & n-k \\ [\mathbf{f}_1^{(1)}]^T & g_{11}^{(1)} \mathbf{e}_1^T \\ F^{(1)}(2:n-k, :) & G^{(1)}(2:n-k, :) \end{pmatrix},$$

where

$$\begin{aligned} \left\| \begin{pmatrix} \mathbf{f}_1^{(1)T} & g_{11}^{(1)} \mathbf{e}_1^T \end{pmatrix} \right\| &= \sigma_1 [(F \ G)], \\ \left\| (F^{(1)}(2:n-k, :) \ G^{(1)}(2:n-k, :)) \right\| &= \sigma_2 [(F \ G)]. \end{aligned}$$

Step 3. Let

$$S = \begin{pmatrix} L & 0 \\ \mathbf{f}_1^{(1)T} & g_{11}^{(1)} \end{pmatrix}.$$

Using a method described in section 3.1, find unit vectors $\bar{\mathbf{u}}_{k+1}, \bar{\mathbf{v}}_{k+1}$ such that

$$S\bar{\mathbf{v}}_{k+1} = \bar{\sigma}_{k+1}\bar{\mathbf{u}}_{k+1},$$

where $\bar{\sigma}_{k+1}$ is the smallest singular value of S .

Step 4. Again, using a method from section 3.2, construct orthogonal matrices $U^{(2)}, V^{(2)} \in \mathbf{R}^{(k+1) \times (k+1)}$ such that

$$\begin{aligned} [U^{(2)}]^T \bar{\mathbf{u}}_{k+1} &= \pm \mathbf{e}_{k+1}, \\ S^{(1)} \leftarrow \begin{pmatrix} L^{(1)} & 0 \\ 0 & \bar{\sigma}_{k+1} \end{pmatrix} &= [U^{(2)}]^T S V^{(2)}, \end{aligned}$$

where $L^{(1)}$ is lower triangular.

$$F^{(2)} \leftarrow (F^{(1)}(2:n-k, :), \ G^{(1)}(2:n-k, 1)) V^{(2)}.$$

If $\bar{\sigma}_{k+1} \geq \epsilon$ (the ϵ -pseudorank is $k+1$), let

$$\bar{L} \leftarrow S^{(1)}, \quad \bar{F} \leftarrow F^{(2)}, \quad \bar{G} \leftarrow G^{(1)}(2:n-k, 2:n-k);$$

else (the ϵ -pseudorank is k)

$$\begin{aligned} \bar{L} &\leftarrow L^{(1)}, \quad \bar{F} \leftarrow \begin{pmatrix} 0 \\ F^{(2)}(:, 1:k) \end{pmatrix}, \\ \bar{G} &\leftarrow \begin{pmatrix} \bar{\sigma}_{k+1} & 0 \\ F^{(2)}(:, k+1) & G^{(2)}(2:n-k, 2:n-k) \end{pmatrix} \end{aligned}$$

end;

$$\begin{aligned} \bar{U} &\leftarrow U \begin{pmatrix} I_k & 0 \\ 0 & U^{(1)} \end{pmatrix} \begin{pmatrix} U^{(2)} & 0 \\ 0 & I_{n-k-1} \end{pmatrix}, \\ \bar{V} &\leftarrow V \begin{pmatrix} I_k & 0 \\ 0 & V^{(1)} \end{pmatrix} \begin{pmatrix} V^{(2)} & 0 \\ 0 & I_{n-k-1} \end{pmatrix}. \end{aligned}$$

The refined decomposition is given by

$$X = \bar{U} \begin{pmatrix} \bar{L} & 0 \\ \bar{F} & \bar{G} \end{pmatrix} \bar{V}^T.$$

The refinement algorithm requires $6(m+n)n + 12nk + c_{iter}n^2$ flops. Thus, this is an $O(mn)$ algorithm for all k .

Algorithm 4.2 does not have a property similar to property (1) of Algorithm 4.1. Fortunately, it does have two properties similar to properties (2) and (3), as follows:

1. $\|(\bar{F} \ \bar{G})\|_F \leq \|(F \ G)\|_F$, where equality occurs if and only if $\|(F \ G)\| = \sigma_{k+1}(C)$.

2. Repeated application forces $\|(\bar{F} \ \bar{G})\|$ to converge to $\sigma_{k+1}(C)$. Under reasonable assumptions, a good approximation of $\sigma_{k+1}(C)$ is produced after one iteration.

The following lemma exactly quantifies the decrease in $\|(F \ G)\|_F$.

LEMMA 4.1. *Assume the terminology of Algorithm 4.2. If $\bar{\sigma}_{k+1} < \epsilon$, then*

$$(4.1) \quad \|(\bar{F} \ \bar{G})\|_F^2 = \|(F \ G)\|_F^2 - \sigma_1^2[(F \ G)] + \sigma_{k+1}^2(S)$$

$$(4.2) \quad = \|(F \ G)\|_F^2 - \left\| \left(\begin{bmatrix} \mathbf{f}_1^{(1)} \\ \vdots \\ \mathbf{f}_1^{(1)} \end{bmatrix}^T, g_{11}^{(1)} \right) \right\|^2 + \bar{\sigma}_{k+1}^2.$$

If $\bar{\sigma}_{k+1} \geq \epsilon$, then

$$(4.3) \quad \|(\bar{F} \ \bar{G})\|_F^2 = \|(F \ G)\|_F^2 - \sigma_1^2[(F \ G)]$$

and

$$\|(\bar{F} \ \bar{G})\| = \sigma_2(F \ G) \leq \|(F \ G)\|.$$

Proof. First, due to orthogonal invariance of the Frobenius norm, we observe that

$$(4.4) \quad \|(F^{(1)} \ G^{(1)})\|_F = \|(F \ G)\|_F,$$

where $\|(F^{(1)} \ G^{(1)})^T \mathbf{e}_1\|$ is the largest singular value of the matrix $(F \ G)$. Second, we observe that

$$(4.5) \quad \left\| \begin{pmatrix} \bar{L} & 0 \\ 0 & \bar{\sigma}_{k+1} \end{pmatrix} \right\|_F = \left\| \begin{pmatrix} L & 0 \\ \mathbf{f}_1^{(1)T} & g_{11}^{(1)} \end{pmatrix} \right\|_F,$$

where $\bar{\sigma}_{k+1}$ is the smallest singular value of S . Thus (4.1) follows from combining (4.4)–(4.5). If $\bar{\sigma}_{k+1} > \epsilon$, then the last row of $S^{(1)}$ is not included in $(\bar{F} \ \bar{G})$, and thus (4.3).

Since $\|(\mathbf{f}_1^{(1)T} \ g_{11}^{(1)})^T\| = \sigma_1[(F \ G)]$ and \mathbf{e}_1 is the left singular vector associated with this singular value, it follows that

$$\|(\bar{F} \ \bar{G})\| = \|(F^{(1)}(2:n-k, :) \ G^{(2)}(2:n-k, :))\| = \sigma_2[(F \ G)]. \quad \square$$

The computation of (4.2) requires $O(n)$ flops.

The following theorem shows that repeated application of Algorithm 4.2 forces $\|(F \ G)\|$ to converge to $\sigma_{k+1}(C)$.

THEOREM 4.2. *Suppose we apply Algorithm 4.2 to $C_0 = \begin{pmatrix} L_0 & 0 \\ F_0 & G_0 \end{pmatrix}$ t times to obtain the matrix $C_t = \begin{pmatrix} L_t & 0 \\ F_t & G_t \end{pmatrix}$. Then*

$$(4.6) \quad \lim_{t \rightarrow \infty} \|(F_t \ G_t)\| = \sigma_{k+1}(C_0).$$

Proof. By using Lemma 4.1 we have

$$(4.7) \quad \|(F_{t+1} \ G_{t+1})\|_F^2 - \|(F_t \ G_t)\|_F^2 = \|(F_t \ G_t)\|^2 - \sigma_{k+1} \left[\left(\begin{bmatrix} L_{t+1} & 0 \\ \mathbf{f}_1^{(t+1)T} & g_{11}^{(t+1)} \end{bmatrix} \right)^2 \right],$$

where $\mathbf{f}_1^{(t+1)T} = F_{t+1}^T \mathbf{e}_1$, $g_{11}^{(t+1)} = \mathbf{e}_1^T G_{t+1} \mathbf{e}_1$.

By the Cauchy interlace theorem [23, pp. 103–104], we have

$$(4.8) \quad \sigma_{k+1} \left[\begin{pmatrix} L_{t+1} & 0 \\ \mathbf{f}_1^{(t+1)T} & g_{11}^{(t+1)} \end{pmatrix} \right] \leq \sigma_{k+1}(C_0) \leq \| (F \ G) \|^2.$$

Thus from (4.7),

$$(4.9) \quad \| (F_{t+1} \ G_{t+1}) \|_F \leq \| (F_t \ G_t) \|_F$$

unless $\| (F_t \ G_t) \| = \sigma_{k+1}(C_0)$. Furthermore, we have

$$(4.10) \quad \sigma_{k+1}(C_0) \leq \| (F_t \ G_t) \|_F \leq \| (F_0 \ G_0) \|_F.$$

Thus by (4.9) and (4.10), $\{ \| (F_t \ G_t) \|_F \}$ defines a monotone, bounded sequence which, in turn, has a limit [18, pp. 47–55]. So we have

$$(4.11) \quad \lim_{t \rightarrow \infty} \| (F_t \ G_t) \|_F = \psi$$

for some $\psi \in \mathbf{R}$.

By (4.7), since the sequence $\{ \| (F_t \ G_t) \|_F \}$ is a Cauchy sequence, and using (4.8), we have

$$\lim_{t \rightarrow \infty} \| (F_t \ G_t) \| = \lim_{t \rightarrow \infty} \sigma_{k+1} \left[\begin{pmatrix} L_t & 0 \\ \mathbf{f}_1^{(t)T} & g_{11}^{(t)} \end{pmatrix} \right] = \sigma_{k+1}(C_0). \quad \square$$

The next theorem shows that just one step of the refinement procedure can yield a value of $\bar{\sigma}_{k+1}$ that is a very good approximation of the $(k+1)$ st singular value of C .

THEOREM 4.3. *For a given matrix C of the form*

$$(4.12) \quad C = \begin{matrix} & k & n-k \\ k & \begin{pmatrix} L & 0 \\ F & G \end{pmatrix} \\ n-k & \end{matrix}$$

Algorithm 4.2 produces a matrix of the form

$$(4.13) \quad \bar{C} = \begin{matrix} & k & n-k \\ k & \begin{pmatrix} \bar{L} & 0 \\ 0 & \bar{\sigma}_{k+1} \mathbf{e}_1^T \\ \tilde{F} & \tilde{G} \end{pmatrix} \\ n-k-1 & \end{matrix}$$

such that

$$(4.14) \quad 0 \leq \sigma_{k+1}(C) - \bar{\sigma}_{k+1} \leq \frac{\sigma_2^2 [(F \ G)]}{\sigma_{k+1}(C) + \bar{\sigma}_{k+1}}.$$

Proof. By orthogonal equivalence, \bar{C} and C have the same singular values. First, we note that

$$\bar{C} \bar{C}^T = \begin{matrix} & k & 1 & n-k-1 \\ k & \begin{pmatrix} \bar{L} \bar{L}^T & 0 & \bar{L} \tilde{F}^T \\ 0 & \bar{\sigma}_{k+1}^2 & \bar{\sigma}_{k+1} \mathbf{e}_1^T \tilde{G}^T \\ \tilde{F} \bar{L}^T & \bar{\sigma}_{k+1} \tilde{G} \mathbf{e}_1 & \tilde{F} \tilde{F}^T + \tilde{G} \tilde{G}^T \end{pmatrix} \\ n-k-1 & \end{matrix}.$$

The Cauchy interlace theorem applied to $\bar{C}\bar{C}^T$ gives us that

$$\bar{\sigma}_{k+1}^2 \leq \sigma_{k+1}^2(\bar{C}) = \sigma_{k+1}^2(C),$$

thereby yielding the upper bound in (4.14).

If we note that

$$\bar{C}^T \bar{C} = \begin{pmatrix} \bar{L}^T \bar{L} & 0 \\ 0 & \bar{\sigma}_{k+1}^2 \mathbf{e}_1 \mathbf{e}_1^T \end{pmatrix} + (\tilde{F} \ \tilde{G})^T (\tilde{F} \ \tilde{G}),$$

then Weyl's monotonicity theorem applied to the $(k+1)$ st eigenvalue of $\bar{C}\bar{C}^T$ is

$$\sigma_{k+1}^2(C) \leq \sigma_{k+1}^2 \left[\begin{pmatrix} \bar{L} & 0 \\ 0 & \bar{\sigma}_{k+1} \end{pmatrix} \right] + \sigma_1^2 [(\tilde{F} \ \tilde{G})].$$

Steps 1–2 of Algorithm 4.2 give us that

$$\sigma_1 [(\tilde{F} \ \tilde{G})] = \sigma_2 [(F \ G)].$$

Steps 3–4 give us that

$$\sigma_{k+1} \left[\begin{pmatrix} \bar{L} & 0 \\ 0 & \bar{\sigma}_{k+1} \end{pmatrix} \right] = \bar{\sigma}_{k+1}.$$

Thus

$$\bar{\sigma}_{k+1}^2 \leq \sigma_{k+1}^2(C) \leq \bar{\sigma}_{k+1}^2 + \sigma_2^2[(F \ G)].$$

Solving for $\sigma_{k+1}(C) - \bar{\sigma}_{k+1}$ yields the upper bound. \square

In the next section, we give some numerical tests that confirm the stated properties of these alternative algorithms.

5. Numerical tests. The following tests were done using MATLAB 6.

Test set. Let $\omega = 1.1\epsilon$, and for a fixed value of k , let $\zeta_k = \ln \omega/k$. Then let

$$D_k = \text{diag}(e^{\zeta_k}, e^{2\zeta_k}, \dots, e^{k\zeta_k}),$$

noting that $\omega = e^{k\zeta_k}$. For $p = \lfloor 100/k \rfloor$, we let

$$\Sigma_k = \text{diag}(D_k, \xi * D_k, \dots, \xi^p D_k, \xi^{p+1} D_k(1 : 100 - p * k)), \quad \xi = 0.9\epsilon.$$

Using a technique for generating random orthogonal matrices given by Stewart [19], we generated $U, V \in \mathbf{R}^{100 \times 100}$ and let

$$X_k = U \Sigma_k V^T.$$

We then computed the QR decomposition

$$X_k = Q C_k^T$$

and used the lower triangular matrix C_k . This is a matrix with exactly k singular values greater than ϵ , but with the singular values clustered.

For each value of $k = 10, 20, \dots, 90$, we generated 10 such matrices. We used values of $\epsilon = 10^{-\ell}$, $\ell = 2, \dots, 8$. The results were similar for all values of k , so we give the results for $\ell = 3$ only.

We computed and plotted the following:

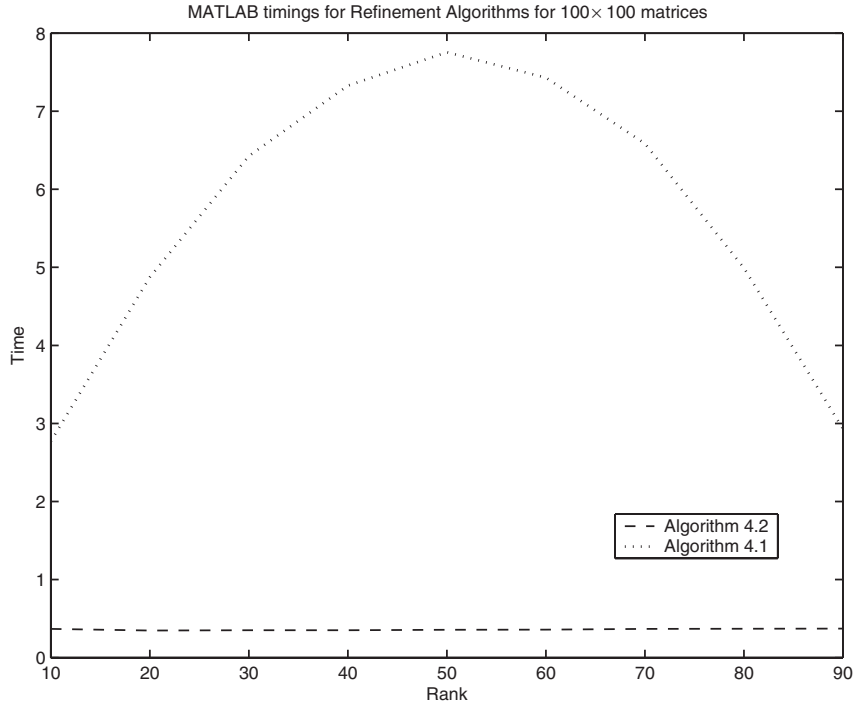


FIG. 5.1. Flop counts of refinement algorithms, $\epsilon = 10^{-3}$.

- MATLAB run times, using the `tic` and `toc` commands (see Figure 5.1).
- Relative error in the k th (in Figure 5.2 (bottom)) and $(k+1)$ st (in Figure 5.2 (top)) singular values. That is, we compute the \log_{10} of

$$\frac{|\sigma_k(C) - \sigma_k(L)|}{\sigma_k(C)}, \quad \frac{|\sigma_{k+1}(C) - \sigma_1(G)|}{\sigma_{k+1}(C)}.$$

- Error in the subspaces. We computed the \log_{10} of

$$|\sin \Theta_R| = \|W_1^T V_2\|_2$$

in Figure 5.1 and the \log_{10} of

$$|\sin \Theta_L| = \|Y_1^T U_2\|_2$$

in Figure 5.3.

The unrefined ULVD was computed using the m-file `hulv_a` from the UTV tools package in [11]. Except for rare cases resulting from inaccuracy of the condition estimators, the ϵ -pseudoranks of these matrices were computed correctly by all of the procedures.

We applied each of Algorithm 4.1 and Algorithm 4.2 after a decomposition computed by `hulv_a`. In Figure 5.1, we plotted the flop counts of the two refinement procedures and the subspace errors, including those from computing a ULVD using `hulv_a` with no refinement at all. In Figure 5.2 we graphed the error in the $(k+1)$ st singular value and the rank error in the k th singular value using `hulv_a`. In Figure 5.3, we graphed the error in the right and left subspaces.

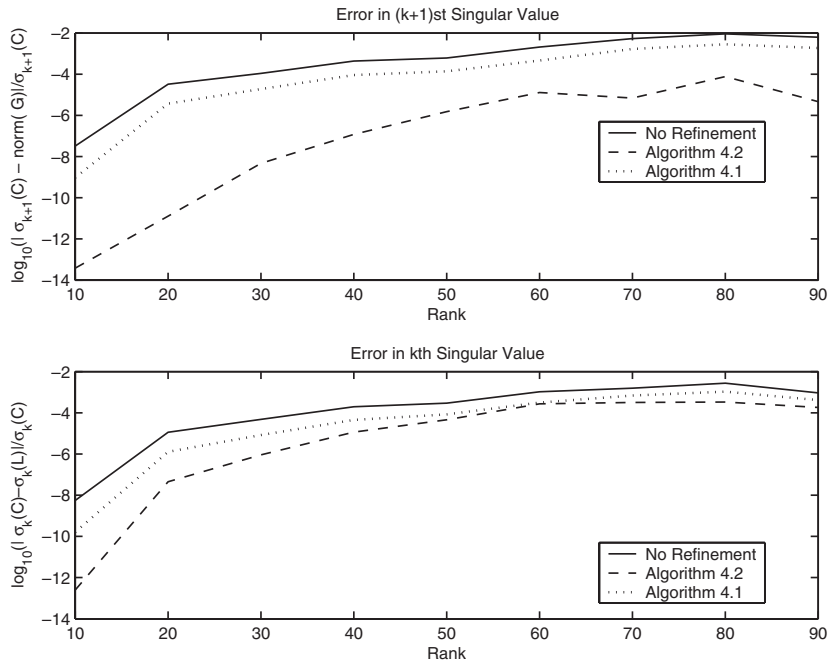


FIG. 5.2. Error in $(k + 1)$ st and k th singular values of refinement algorithms, $\epsilon = 10^{-3}$.




FIG. 5.3. Error in right and left subspaces of refinement algorithms, $\epsilon = 10^{-3}$.

As one would predict, Algorithm 4.2 is much faster than Algorithm 4.1. Algorithm 4.2 significantly improves the estimate of $\sigma_{k+1}(C)$ given by $\sigma_1(G)$ and slightly improves the estimate of $\sigma_k(C)$ given by $\sigma_k(L)$. It does both better than Algorithm 4.1. It also makes the angles Θ_R and Θ_L with the right and left singular subspaces small and does so as well as or better than Algorithm 4.1.

6. Conclusion. The ULVD has already been shown to be a good substitute for the SVD in circumstances where frequent updates and downdates are done.

We propose a refinement procedure that requires $O(mn)$ flops that is just as effective as the $O(mn^2)$ flops procedures proposed in [17]. Our tests above show that this procedure successfully improves both the singular value and subspace estimates.

REFERENCES

- [1] J. L. BARLOW, *Modification and maintenance of ULV decompositions*, in Applied Mathematics and Scientific Computing, Z. Drmač et al., eds., Kluwer Academic, New York, 2003, pp. 31–63.
- [2] J. L. BARLOW, H. ERBAY, AND I. SLAPNIČAR, *An Alternative Algorithm for Refinement of ULV Decompositions*, Tech. report, Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA, 2004; available online from http://www.cse.psu.edu/~barlow/ULV_Hasan3.ps.
- [3] J. L. BARLOW AND P. A. YOON, *Solving recursive TLS problems using the rank-revealing ULV decomposition*, in Recent Advances in Total Least Squares Techniques and Errors-in-Variables Modeling, S. Van Huffel, ed., SIAM, Philadelphia, PA, 1997, pp. 117–126.
- [4] J. L. BARLOW, P. A. YOON, AND H. ZHA, *An algorithm and a stability theory for downdating the ULV decomposition*, BIT, 36 (1996), pp. 14–40.
- [5] A. BJÖRCK, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, PA, 1996.
- [6] P. A. BUSINGER AND G. H. GOLUB, *Linear least squares solutions by Householder transformations*, Numer. Math., 7 (1965), pp. 269–278.
- [7] C. DAVIS AND W. M. KAHAN, *The rotation of eigenvectors by a perturbation*. III, SIAM J. Numer. Anal., 7 (1970), pp. 1–46.
- [8] H. ERBAY, J. L. BARLOW, AND Z. ZHANG, *A modified Gram–Schmidt–based downdating technique for ULV decompositions with applications to recursive TLS problems*, Comput. Statist. Data Anal., 41 (2003), pp. 195–211.
- [9] D. K. FADDEEV, V. N. KUBLANOVSKAYA, AND V. N. FADDEEVA, *Sur les Systèmes Linéaires Algébriques de Matrices Rectangulaires et Mal-Conditionnées*, Vol. VII, Editions Centre Nat. Recherche Sci., Paris, 1968, pp. 161–170.
- [10] R. D. FIERRO AND J. R. BUNCH, *Bounding the subspaces from rank revealing two-sided orthogonal decompositions*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 743–759.
- [11] R. FIERRO, P. C. HANSEN, AND P. S. KIRK, *UTV tools: Matlab templates for rank-revealing UTV decompositions*, Numer. Algorithms, 20 (1999), pp. 165–194.
- [12] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, Baltimore, MD, 1996.
- [13] M. GU AND S. C. EISENSTAT, *Downdating the singular value decomposition*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 793–810.
- [14] R. J. HANSON AND C. L. LAWSON, *Extensions and applications of the Householder algorithm for solving linear least squares problems*, Math. Comp., 23 (1969), pp. 787–812.
- [15] I. T. JOLIFFE, *Principal Component Analysis*, Springer-Verlag, New York, 1986.
- [16] C. L. LAWSON AND R. J. HANSON, *Solving Least Squares Problems*, Prentice-Hall, Englewood Cliffs, NJ, 1974.
- [17] R. MATHIAS AND G. W. STEWART, *A block QR algorithm and the singular value decomposition*, Linear Algebra Appl., 182 (1993), pp. 91–100.
- [18] W. RUDIN, *Principals of Mathematical Analysis*, 3rd ed., McGraw-Hill, New York, 1976.
- [19] G. W. STEWART, *The efficient generation of random orthogonal matrices with an application to condition estimators*, SIAM J. Numer. Anal., 17 (1980), pp. 403–409.
- [20] G. W. STEWART, *An updating algorithm for subspace tracking*, IEEE Trans. Signal Process., 40 (1992), pp. 1535–1541.
- [21] G. W. STEWART, *Updating a rank-revealing ULV decomposition*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 494–499.

- [22] S. VAN HUFFEL AND J. VANDEWALLE, *The Total Least Squares Problem: Computational Aspects and Analysis*, SIAM, Philadelphia, PA, 1991.
- [23] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford University Press, London, 1965.
- [24] P. A. YOON AND J. L. BARLOW, *An efficient rank detection procedure for modifying the ULV decomposition*, BIT, 38 (1998), pp. 781–801.