UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

On the Convergence of Stewart's QLP Algorithm for Approximating the SVD

David A. Huckaby
Tony F. Chan

May 2002
CAM Report 02-29

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90095-1555

http://www.math.ucla.edu/applied/cam/index.html

On the Convergence of Stewart's QLP Algorithm for Approximating the SVD

David A. Huckaby and Tony F. Chan

May 22, 2002

Abstract

The pivoted QLP decomposition, introduced by G. W. Stewart [14], represents the first two steps in an algorithm which approximates the SVD. The matrix $A\Pi_0$ is first factored as $A\Pi_0 = QR$, and then the matrix $R^T\Pi_1$ is factored as $R^T\Pi_1 = PL^T$, resulting in $A = Q\Pi_1 LP^T\Pi_0^T$, with Q and P orthogonal, L lower-triangular, and Π_0 and Π_1 permutation matrices. Stewart noted that the diagonal elements of L approximate the singular values of A with surprising accuracy. In this paper, we provide mathematical justification for this phenomenon. If there is a gap between σ_k and σ_{k+1}, partition the matrix L into diagonal blocks L_{11} and L_{22} and off-diagonal block L_{21}, where L_{11} is k-by-k. We show that the convergence of $(\sigma_j(L_{11})^{-1} - \sigma_j^{-1})/\sigma_j^{-1}$ for $j = 1, \ldots, k$, and of $(\sigma_j(L_{22}) - \sigma_{k+j})/\sigma_{k+j}$, for $j = 1, \ldots, n-k$ are all quadratic in the gap ratio σ_{k+1}/σ_k. The worst case is therefore at the gap, where the absolute errors $\|L_{11}^{-1}\| - \sigma_k^{-1}$ and $\|L_{22}\| - \sigma_{k+1}$ are thus cubic in σ_k^{-1} and σ_{k+1}, respectively. One order of covergence is due to the rank-revealing pivoting in the first step; then, because of the pivoting in the first step, two more orders are achieved in the second step. Our analysis assumes that $\Pi_1 = I$, that is, that pivoting is done only on the first step. The algorithm can be continued beyond the first two steps, and we make some observations concerning the asymptotic convergence. For example, we point out that repeated singular values can accelerate convergence of individual elements. This, in addition to the relative convergence to all of the singular values being quadratic in the gap ratio, indicates that the QLP decomposition can be powerful even when the ratios between neighboring singular values are close to one.

1 Introduction

1.1 Overview

One of the most important tools for analyzing a matrix A is the Singular Value Decomposition (SVD): If A is a real m-by-n matrix, then there exist orthogonal matrices

$$U = [u_1, \dots, u_m] \in \mathbb{R}^{m \times m} \text{ and } V = [v_1, \dots, v_n] \in \mathbb{R}^{n \times n}$$

such that

$$U^T AV = diag(\sigma_1, \dots, \sigma_p) \equiv \Sigma \in \mathbb{R}^{m \times n} \; p = \min\{m, n\},$$

where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p \geq 0$ [6, Theorem 2.5.2]. The vectors u_i and v_i are called, respectively, the ith *left singular vector* and ith *right singular vector* of A. The σ_i are called the *singular values* of A.

While the SVD gives a lot of information about A, it is rather costly to compute. A worthy goal is to develop a method that provides information close to the quality that the SVD provides but which runs much faster.

One attempt at gaining SVD-type information is, of course, the QR factorization [6, Section 5.2]: Let $A \in \mathbb{R}^{m \times n}$ have rank n. Then A can be written uniquely in the form $A = QR$, where Q is an $m \times n$ orthogonal matrix and R is an $n \times n$ upper triangular matrix with positive diagonal elements. This can be made into a full blown decomposition which handles rank-deficient A. For details, see [13, Chap. 4, Sec. 1].

The QR factorization is very cheap to compute, relative to the SVD. Since Q is orthogonal, we know that $\|A\|_2 = \|R\|_2$, so that the singular values of R are the same as those of A. As we have effectively reduced a dense matrix A to an upper triangular matrix R, we might hope that the diagonal elements of R are approximations to the singular values of A.

This is clearly not true in general. For example, if the norm of the first column of A is small compared to the norms of the rest of the columns, then r_{11} is small and comes nowhere close to $\sigma_1 = \|A\|_2$. All that is readily apparent is that $r_{11} \leq \sigma_1$ and $r_{nn} \geq \sigma_n$.

We can introduce column pivoting into the computation of the QR factorization [7], and this makes a huge difference. Pivoting ensures that the diagonal elements in the computed R-factor are in sorted order, and these are often rough approximations of the corresponding singular values. They can also be used for gap revelation, which of course requires less accuracy. Pivoting also provides at least one bound that its unpivoted brother could not promise. The fact that r_{11} equals the norm of the largest column of A gives rise to the bound $r_{11} \geq \sigma_1/\sqrt{n}$ (see page 14).

Note that although the column pivoted QR factorization provides this useful bound for the error in approximating σ_1 by r_{11}, only an exponential bound exists for σ_n and r_{nn} [4]. So although in practice the diagonal elements of R are rough approximations to the singular values, this need not always be the case. (Indeed, the Kahan matrix provides a well-known counterexample. [8])

This is one consideration that motivates the *rank revealing QR factorization* (RRQR) [5, 1]. If R is partitioned as

$$R = \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix},$$

an RRQR algorithm tries to maximize the smallest singular value of R_{11} and/or minimize the largest singular value of R_{22} [2]. This essentially means making, respectively, $\|R_{11}\|$ as large as possible and $\|R_{22}\|$ as small as possible. From the interlacing property of singular values, $\sigma_{min}(R_{11}) \leq \sigma_k(A)$ and $\sigma_{max}(R_{22}) \geq \sigma_{k+1}(A)$. So an RRQR factorization provides bounds on the singular values of A in terms of the norms of the blocks.

In terms of tracking the singular values, the RRQR algorithms tend to perform about as well as the pivoted QR factorizations, and they come with guarantees. Whereas the pivoted QR factorization can completely fail (like on Kahan's example [8]), an RRQR algorithm is guaranteed to work within the bounds it provides. So for example, one algorithm which tries to minimize $\|R_{22}\|$ can promise that $r_{nn} \leq \sqrt{n}\,\sigma_n$, a bound that the column pivoted QR factorization cannot provide.

Stewart [14] introduced another candidate for SVD-quality information with minimal cost having the QR factorization as its only building block. The pivoted QLP decomposition requires only the work of two QR factorizations, and one of them need not even be pivoted. Yet despite its simplicity and speed, the decomposition provides approximations to the singular subspaces of A and gives excellent approximations to *all* of the singular values of A. It is the purpose of this paper to provide some theoretical underpinning for this observation.

The paper is organized as follows. In § 2 we discuss the QLP decomposition and illustrate just how good it is at tracking the singular values. We also consider taking it beyond just two successive QR factorizations, which leads to an algorithm that asymptotically calculates the SVD. In discussing the convergence of this algorithm in the absence of pivoting, we point out the connection to the QR algorithm and mention past work done by Mathias and Stewart [9] and Chandrasekaran and Ipsen [3]. In § 3 we study the convergence of the QLP decomposition. The decomposition is obtained by performing a pivoted QR factorization of A as $A = QR$, where we have incorporated the pivoting into the matrix Q, and then a pivoted QR factorization of R^T as $R^T = PL$. If there is a gap between σ_k and σ_{k+1}, partition the matrix L into diagonal blocks L_{11} and L_{22} and off-diagonal block L_{21}, where L_{11} is k-by-k. We show that the convergence of $(\sigma_j(L_{11})^{-1} - \sigma_j^{-1})/\sigma_j^{-1}$ for $j = 1, \ldots, k$, and of $(\sigma_j(L_{22}) - \sigma_{k+j})/\sigma_{k+j}$, for $j = 1, \ldots, n-k$ are all quadratic in the gap ratio σ_{k+1}/σ_k. The worst case is therefore at the gap, where the absolute errors $\|L_{11}^{-1}\| - \sigma_k^{-1}$ and $\|L_{22}\| - \sigma_{k+1}$ are thus cubic in σ_k^{-1} and σ_{k+1}, respectively. It turns out that one order is due to the rank-revealing pivoting in the first step; then, the pivoting has provided a springboard so that two more orders are achieved in the second step. Our analysis assumes that $\Pi_1 = I$, that is, that pivoting is done only on the first step. That we have relative convergence quadratic in the gap ratio for *all* of the singular values helps show why the QLP decomposition can work so well even when most of ratios between singular values are close to one. It only takes one significant gap for the decomposition to work very well. We provide numerical experiments to illustrate the results.

In § 4 we present some observations concerning the asymptotic convergence of the algorithm. The off-diagonal elements converge in interesting patterns, and from them we can discern the asymptotic rates of convergence of the diagonal elements. From this perspective of individual element convergence, we again see the phenomenon of repeated singular values (or singular values that are close) accelerating convergence. The results are obtained by analyzing the actions of using Givens rotations in computing the QR factorizations, and numerical experiments are again provided to illustrate the theory.

2 QLP: An Approximate SVD

2.1 The Pivoted QLP Decomposition

The QLP decomposition was introduced by G. W. Stewart [14], who observed its potency in rank revelation, singular value approximation, and gap revelation. Let us call the diagonal elements in the R matrix of a QR factorization of A the R-values of A. Noting that the R-values are rough approximations of the singular values, Stewart suggested taking the pivoted QR factorization and then triangularizing on the right, obtaining the factorization $A = Q\Pi_1 LP^T \Pi_0^T$. If we include the permutation matrices Π_1 and Π_0^T as part of A and R^T, we have $A = QLP^T$, called the pivoted QLP decomposition of A. Note that the second step is equivalent to performing a QR factorization on R^T, obtaining $R^T = PL^T$. Also note that L is lower-triangular. So the decomposition amounts to taking two pivoted QR factorizations and thus factorizing A into the product of an orthogonal matrix, a lower-triangular matrix, and another orthogonal matrix. The diagonal elements of L are called the L-values of A.

Stewart showed empirically that the L-values track the singular values surprisingly well—far better than the R-values. See Figure 1.

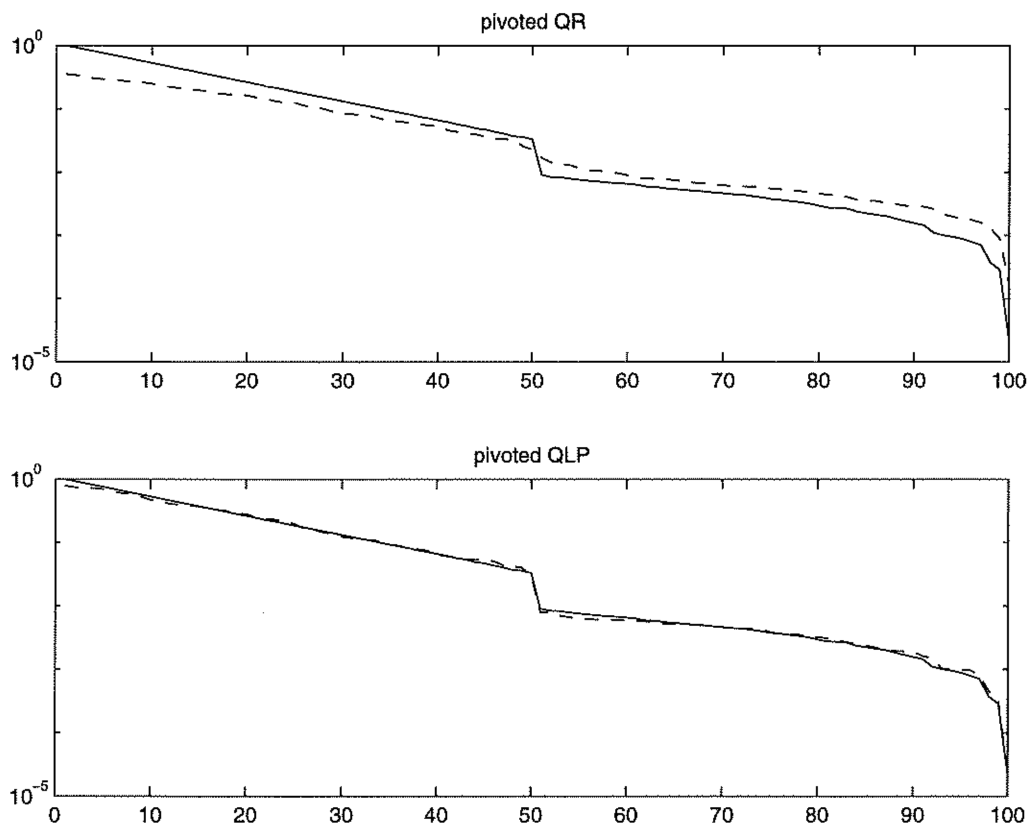

Figure 1: QR vs QLP. Here the solid lines are the singular values of a 100-by-100 matrix that has a gap between σ_{50} and σ_{51}. The dotted line represents the R-values in the first plot and the L-values in the second.

Note from Figure 1 that not only do the L-values identify the gap far better than the R-values, they also approximate the singular values. So with only the extra cost of one more QR factorization, we get very good information—almost SVD-quality information in many situations. (Note that the QLP decomposition is a special case of the ULV decomposition, also introduced by Stewart [11, 12]. The usual ULV and URV decompositions are also rank-revealing but do not attempt to approximate the singular values. Their main selling point is that they are easily updated.)

2.2 The QLP Iteration

Stewart points out that without pivoting, the decomposition represents the first two steps in an iterative algorithm that actually computes the SVD [9]. Let us call this iterative algorithm the *QLP iteration*. In each step after the first, we just compute the QR factorization of the transpose of the R factor produced by the last step. Here is the algorithm:

1. Compute the QR factorization of A, obtaining $A = Q_0 R_0$.
2. Compute the QR factorization of R_0^T, obtaining $R_0^T = Q_1 R_1$.
3. Compute the QR factorization of R_1^T, obtaining $R_1^T = Q_2 R_2$.
4. Continue in the same way.

Notice that if we stop after the second step and perform the QR factorizations in steps 1 and 2 with pivoting, then this is just the QLP decomposition. If on the other hand we do not pivot and do not stop after the second step but continue on, we obtain something akin to the QR algorithm for computing eigenvalues and eigenvectors, which follows [15]:

$$
\begin{array}{ll}
\multicolumn{2}{l}{A^{(0)} = A} \\
\multicolumn{2}{l}{\textbf{for}\ \ k = 1, 2, \ldots} \\
\quad Q^{(k)} R^{(k)} = A^{(k-1)} & \text{QR factorization of } A^{(k-1)} \\
\quad A^{(k)} = R^{(k)} Q^{(k)} & \text{Switch factors}
\end{array}
$$

The $A^{(k)}$ converge to a diagonal matrix whose elements are the eigenvalues of A in decreasing order.

To see how the QLP iteration relates to the QR algorithm, define the matrix $A^{(i)} = R_{2i}^T R_{2i}$. Now from the QLP iteration, $R_{2i}^T = Q_{2i+1}R_{2i+1}$, so that $A^{(i)} = Q_{2i+1}R_{2i+1}R_{2i}$. Since the product of two upper-triangular matrices is upper-triangular, this equation represents a QR factorization of $A^{(i)}$, the first step of the QR algorithm. To perform the second step, we switch the two factors [9]:

$$
\begin{aligned}
A^{(i+1)} &= R_{2i+1}R_{2i}Q_{2i+1} \\
&= R_{2i+1}R_{2i+1}^T \\
&= R_{2i+1}Q_{2i+2}^T Q_{2i+2}R_{2i+1}^T \\
&= R_{2i+2}^T R_{2i+2}.
\end{aligned}
$$

We see that every two steps of the QLP iteration (excluding the first) on the R_i are equivalent to one step of the QR algorithm on the $R_{2i}^T R_{2i}$. Since the $A^{(i)}$ converge to a diagonal matrix whose elements are the eigenvalues of $A^{(0)} = R_0^T R_0$ in decreasing order, the R_{2i} converge to a diagonal matrix whose elements are the singular values of R_0 in decreasing order. (A similar argument can be used for the R_i when i is odd.)

Some convergence results for the unpivoted version of the iteration were given by Mathias and Stewart [9] and Chandrasekaran and Ipsen [3]. Let R be the upper-triangular matrix at one step of the iteration and let R' be the upper-triangular matrix at the next step. Partition the n-by-n matrices as

$$
R = \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix} \quad \text{and} \quad R' = \begin{pmatrix} R'_{11} & R'_{12} \\ 0 & R'_{22} \end{pmatrix},
$$

where R_{11} and R'_{11} are k-by-k. Mathias and Stewart showed that if $\rho = \|R_{22}\|/\inf(R_{11}) < 1$, then $\|R'_{12}\| \leq \rho\|R_{12}\|$. Chandrasekaran and Ipsen studied not only the convergence of the R matrices, but also the convergence of the singular vectors of the R matrices as well. They provided some monotonic convergence results and some asymptotic convergence results and also suggested preceeding the iteration with an RRQR algorithm to make ρ as small as possible.

We would like to point out that the QLP iteration also computes the singular vectors. At each iteration, an orthogonal matrix is computed as part of the QR factorization. A product of these accumulates on each side of the $R^{(i)}$, the product on the left converging to U and the product on the right converging to V^T. In the case of the QLP decomposition, that is, stopping after two iterations, there is only one matrix on each side, Q and P. Error bounds on how well they approximate the singular subspaces are available. See [9] and [13, Chap. 5, Thm. 2.3]. Our focus will be on the convergence of the triangular matrix to the singular values.

3 Convergence of the QLP Decomposition

3.1 The Extreme Singular Values

First we will see how well the QLP decomposition approximates the extremal singular values, σ_1 and σ_n. In studying the QLP decomposition, we let the notation reflect the fact that we view the decomposition as the first two steps of the QLP iteration. Thus we use $Q^{(0)}$ instead of Q, $Q^{(1)}$ instead of P, and $(R^{(1)})^T$ instead of L. Stewart pointed out that the pivoting in the first step was crucial, but necessary in the second step only to avoid "certain contrived counterexamples." So to simplify the analysis, we will assume no pivoting on the second step. This is fine, since the results can only get better with pivoting.

We are not here seeking precise bounds, but only rough estimates sufficient to explain the good convergence after only two steps. In analyzing the first step, we note that all rank-revealing-type algorithms perform roughly equally most of the time. This frees us to use bounds from any RRQR algorithm in our derivations, while still tacitly assuming a pivoted QR decomposition in the first step. We will use the bounds provided for "Hybrid-III" by Chandrasekaran and Ipsen [2].

Since we are interested in bounding the error in approximating σ_n, we let $k = n-1$. In this case, the relevant bounds are:

$$|r_{nn}^{(0)}| \leq \sqrt{n}\,\sigma_n, \tag{1}$$

$$\inf(R_{11}^{(0)}) \geq \frac{1}{\sqrt{2(n-1)}}\sigma_n. \tag{2}$$

(The bound (1) is originally due to Golub and Van Loan [6].)

After the first step, we no longer have a full matrix A but an upper-triangular matrix $R^{(0)}$, so we may employ the analysis of Mathias and Stewart [9]. Their paper is divided into two parts. In the first, they studied the unpivoted iteration on block upper-triangular matrices. They showed [9, Theorem 2.1] that under the assumption

$$\rho^{(i)} = \frac{\|R_{22}^{(i)}\|}{\inf(R_{11}^{(i)})} < 1, \tag{3}$$

the following hold:

$$\|R_{12}^{(i+1)}\| \leq \rho^{(i)}\|R_{12}^{(i)}\|, \tag{4}$$

$$\sigma_j(R_{22}^{(i+1)}) \leq \sigma_j(R_{22}^{(i)}), \qquad j = 1, \ldots, n-k, \tag{5}$$

$$\sigma_j(R_{11}^{(i+1)}) \geq \sigma_j(R_{11}^{(i)}), \qquad j = 1, \ldots, k. \tag{6}$$

In the second part of their paper, they used these results to approximate the singular values of a block upper-triangular matrix R by the singular values of a block diagonal matrix $\hat{R}$. Partition R and $\hat{R}$ as

$$R = \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix} \quad \text{and} \quad \hat{R} = \begin{pmatrix} R_{11} & 0 \\ 0 & R_{22} \end{pmatrix}.$$

In general, the singular values of $\hat{R}$ differ from those of R by no more than $\|R_{12}\|$ [6, Cor. 8.6.2]. Mathias and Stewart show that if

$$\rho = \frac{\|R_{22}\|}{\inf(R_{11})} < 1, \tag{7}$$

then

$$\sigma_j(R)/\sigma_j(\hat{R}) = 1 + \mathcal{O}\left(\frac{\|R_{12}\|^2}{(1-\rho^2)[\inf(R_{11})]^2}\right).$$

We will assume that (3) holds with $i = 0$, that is, after the first step. This is not an unreasonable assumption, since the pivoted QR factorization of the first step roughly orders the embryonic singular values. Note that if A has j zero singular values, then this initial pivoting will put zeros as the last j elements on the diagonal of $R^{(0)}$. The matrix can then be deflated (the last j rows and j columns of $R^{(0)}$ discarded), and the iteration continued with $R^{(0)}$ having all non-zero singular values. We may therefore assume that there are no zero singular values.

Now we are ready to state our first theorem.

Theorem 3.1 *Let A be an m-by-n matrix and let $\sigma_{n-1}(A) > \sigma_n(A)$. Let $R^{(0)}$ be the R-factor in the pivoted QR factorization of A, $A\Pi = Q^{(0)}R^{(0)}$ and let $R^{(1)}$ be the R-factor in the unpivoted QR factorization of $(R^{(0)})^T$, $(R^{(0)})^T = Q^{(1)}R^{(1)}$. Assume that the bounds (1) and (2) hold and that $\rho^{(0)} = |r_{nn}^{(0)}|/inf(R_{11}^{(0)}) < 1$.*

Then

$$|r_{nn}^{(1)}| - \sigma_n \leq \frac{\sigma_n^3}{\sigma_{n-1}^2} \mathcal{O}\left(\frac{n^{\frac{5}{2}} \|(R_{12}^{(0)})\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2} \right). \tag{8}$$

Proof:

Since $\rho^{(0)} = |r_{nn}^{(0)}|/\inf(R_{11}^{(0)}) < 1$, by (5) and (6) we have $|r_{nn}^{(1)}| \leq |r_{nn}^{(0)}|$ and $\|R_{11}^{(1)}\| \geq \|R_{11}^{(0)}\|$, so that $\rho^{(1)} = |r_{nn}^{(1)}|/\inf(R_{11}^{(1)}) < 1$ also. Hence Mathias and Stewart's result [9, Theorem 3.1] gives

$$\frac{\sigma_n}{|r_{nn}^{(1)}|} \geq 1 - \mathcal{O}\left(\frac{\|R_{12}^{(1)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2} \right),$$

or

$$|r_{nn}^{(1)}| - \sigma_n \leq |r_{nn}^{(1)}| \mathcal{O}\left(\frac{\|R_{12}^{(1)}\|^2}{[1-(\rho^{(1)})^2)(\inf(R_{11}^{(1)})]^2} \right),$$

From this, the result follows:

$$
\begin{aligned}
|r_{nn}^{(1)}| - \sigma_n &\le |r_{nn}^{(1)}| \mathcal{O}\left(\frac{\|R_{12}^{(1)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right) \\
&\le |r_{nn}^{(1)}| \mathcal{O}\left(\frac{(\rho^{(0)})^2\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right) \qquad (9)\\
&\le |r_{nn}^{(0)}| \mathcal{O}\left(\frac{(\rho^{(0)})^2\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right) \qquad (10)\\
&\le \frac{|r_{nn}^{(0)}|^3}{\inf(R_{11}^{(0)})^2} \mathcal{O}\left(\frac{\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right) \\
&\le \frac{\sigma_n^3}{\sigma_{n-1}^2} \mathcal{O}\left(\frac{n^{\frac{5}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right). \qquad (11)
\end{aligned}
$$

The inequality (9) follows from (4), (10) follows from (5), and (11) follows from (1) and (2). □

Note that here $R_{12}^{(1)}$ is a vector. Theorem 3.1 gives a bound on the error in approximating the smallest singular value σ_n of A by $|r_{nn}^{(1)}|$, the absolute value of the final element on the diagonal of $R^{(1)}$. We see that the error $|r_{nn}^{(1)}| - \sigma_n$ is cubic in σ_n. From equation (10), we see that one order comes from the $|r_{nn}^{(0)}|$ factor in front, which is bounded in the rank-revealing first step while two orders come from the $(\rho^{(0)})^2$ factor contributed in the second step. The inequality also shows that the relative error $(|r_{nn}^{(1)}| - \sigma_n)/\sigma_n$ is quadratic in σ_n/σ_{n-1}.

We illustrate Theorem 3.1 on a 30-by-30 matrix. We fix the largest 29 singular values to be spaced evenly from 1 to 10. We perform the QLP decomposition (with pivoting on the first step and no pivoting on the second step) five times, allowing the smallest singular value, σ_{30}, to take on the values $10^{-1}, 10^{-2}, \ldots, 10^{-5}$. We expect the absolute error $|r_{30,30}^{(1)}| - \sigma_{30}$ to decrease by 10^{-3} on each run and the relative error $(|r_{30,30}^{(1)}| - \sigma_{30})/\sigma_{30}$ to decrease by 10^{-2} on each run. This is verified in Figure 2.

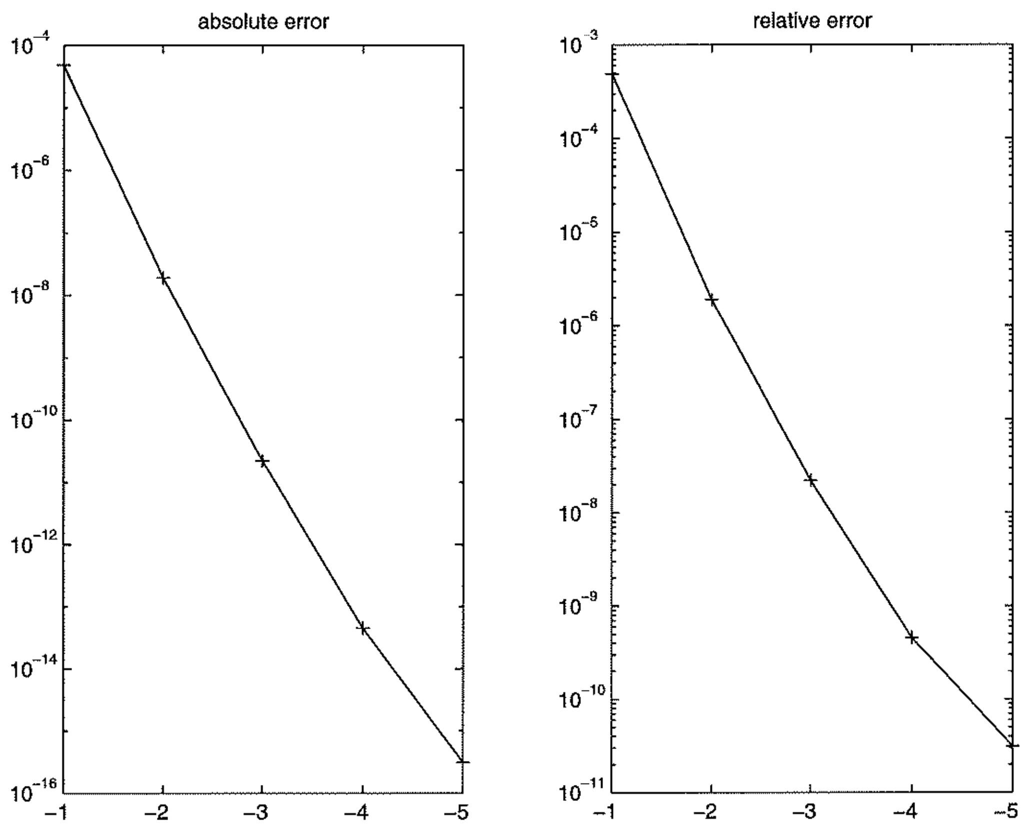

Figure 2: As σ_{30} decreases from 10^{-1} down to 10^{-5}, the absolute error $|r_{nn}^{(1)}| - \sigma_n$ decreases as the cube of σ_{30}, and the relative error $(|r_{nn}^{(1)}| - \sigma_n)/\sigma_n$ decreases as the square of $\sigma_{30}/\sigma_{29} = \sigma_{30}/1 = \sigma_{30}$.

There is an analogous result for $r_{11}^{(1)}$. Because the pivoted QR factorization makes $|r_{11}^{(0)}|$ equal to the norm of the largest column of A, $|r_{11}^{(0)}| = \max_{1\leq j\leq n} \|Ae_j\|_2^2$, where e_j is the j-th canonical vector, we do not need to borrow an RRQR bound for $|r_{11}^{(0)}|$. For any $A \in \mathbb{R}^{l\times m}$, we have

$$\max_{1\leq j\leq n} \|Ae_j\|_2 \geq \frac{\|A\|_2}{\sqrt{n}}. \tag{12}$$

This is a standard result: $\|A\|_2^2 = \max_{\|x\|_2=1} \|Ax\|_2^2 \leq \|A\|_F^2 \leq n \max_{1\leq j\leq n} \|Ae_j\|_2^2$.

The first inequality is from, for example, Theorem 2.5 on p. 175 of [10].

When we take the column pivoted QR factorization of a matrix A, the absolute value of the element r_{11} in the upper-left corner is equal to $\max_{1\le j\le n}\|Ae_j\|_2$. Equation (12) therefore tells us that $|r_{11}^{(0)}|$ is greater than or equal to $\sigma_1/\sqrt{n}$. So here we have an RRRQ bound automatically. The one RRQR bound assumed for this case is [2]:

$$\|R_{22}^{(0)}\| \le \sqrt{2(n-1)}\,\sigma_2. \tag{13}$$

The theorem is very similar to Theorem 3.1.

Theorem 3.2 *Let A be an m-by-n matrix and let $\sigma_1(A) > \sigma_2(A)$. Let $R^{(0)}$ be the R-factor in the pivoted QR factorization of A, $A\Pi = Q^{(0)}R^{(0)}$ and let $R^{(1)}$ be the R-factor in the unpivoted QR factorization of $(R^{(0)})^T$, $(R^{(0)})^T = Q^{(1)}R^{(1)}$. Assume that the bound (13) holds and that $\rho^{(0)} = \|R_{22}^{(0)}\|/|r_{11}^{(0)}| < 1$.*

Then

$$|r_{11}^{(1)}|^{-1} - \sigma_1^{-1} \le \frac{\sigma_2^2}{\sigma_1^3}\mathcal{O}\left(\frac{n^{\frac{5}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right). \tag{14}$$

Proof:

Since $\rho^{(0)} = \|R_{22}^{(0)}\|/|r_{11}^{(0)}| < 1$, by (5) and (6) we have $\|R_{22}^{(1)}\| \le \|R_{22}^{(0)}\|$ and $|r_{11}^{(1)}| \ge |r_{11}^{(0)}|$, so that $\rho^{(1)} = \|R_{22}^{(1)}\|/|r_{11}^{(1)}| < 1$. Hence Mathias and Stewart's result [9, Theorem 3.1] gives

$$\frac{|r_{11}^{(1)}|}{\sigma_1} \ge 1 - \mathcal{O}\left(\frac{\|R_{12}^{(1)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right).$$

or

$$|r_{11}^{(1)}|^{-1} - \sigma_1^{-1} \leq |r_{11}^{(1)}|^{-1} \mathcal{O}\left(\frac{\|R_{12}^{(1)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right).$$

From this, the result follows:

$$\begin{aligned}
|r_{11}^{(1)}|^{-1} - \sigma_1^{-1} &\leq |r_{11}^{(1)}|^{-1} \mathcal{O}\left(\frac{\|R_{12}^{(1)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right) \\
&\leq |r_{11}^{(1)}|^{-1} \mathcal{O}\left(\frac{(\rho^{(0)})^2\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right) && (15) \\
&\leq |r_{11}^{(0)}|^{-1} \mathcal{O}\left(\frac{(\rho^{(0)})^2\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right) && (16) \\
&\leq \frac{\|R_{22}^{(0)}\|^2}{|r_{11}^{(0)}|^3} \mathcal{O}\left(\frac{\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right) \\
&\leq \frac{\|R_{22}^{(0)}\|^2}{\sigma_1^3} \mathcal{O}\left(\frac{n^{\frac{3}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right) && (17) \\
&\leq \frac{\sigma_2^2}{\sigma_1^3} \mathcal{O}\left(\frac{n^{\frac{5}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2](r_{11}^{(1)})^2}\right). && (18)
\end{aligned}$$

The inequality (15) follows from (4), (16) follows from 6, (17) follows from equation (12), and (18) follows from 13. □

Note that in Theorem 3.2, the quantities in question are reciprocals, and the convergence is cubic in $1/\sigma_1$. We illustrate the theorem using a 30-by-30 matrix as we did Theorem 3.2. This time we fix the smallest 29 singular values to be spaced evenly from 0.1 to 1. We allow σ_1 to increase from 10 to 10^5 over five runs, and we expect to see the absolute error decrease by a factor of 10^3 and the relative error decrease by a factor of 10^2 with each run. See Figure 3 for the results.

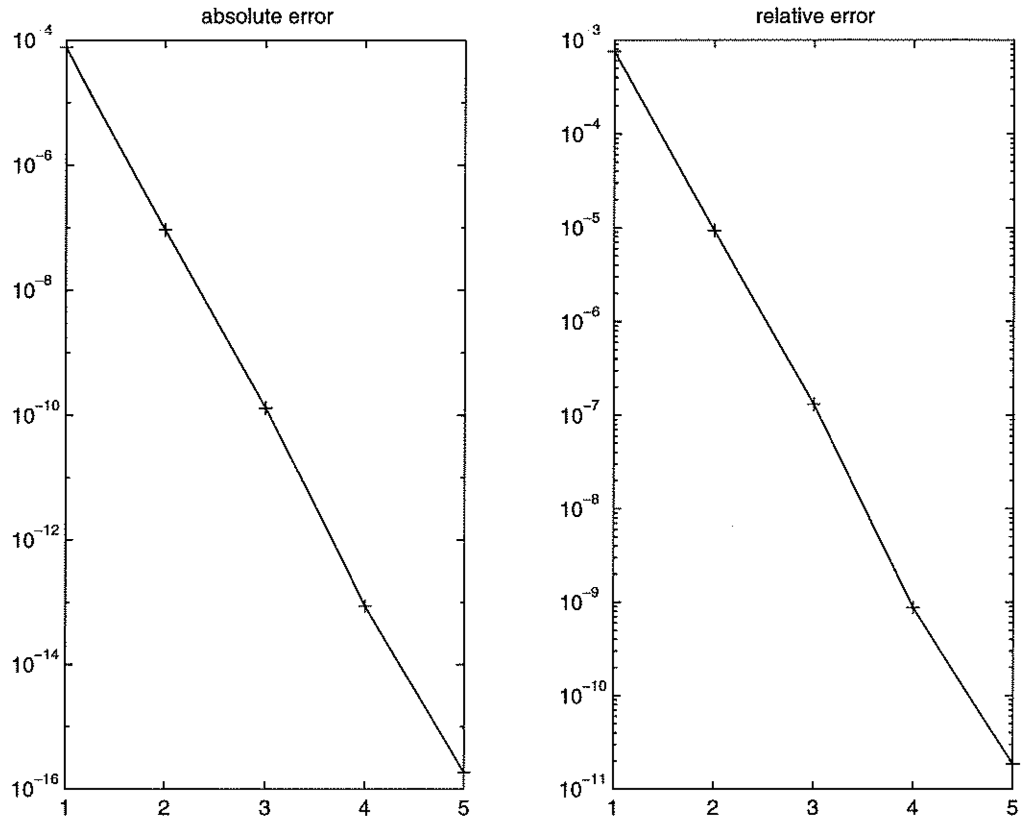

Figure 3: As σ_1 increases from 10^1 up to 10^5, the absolute error $|r_{11}^{(1)}|^{-1} - \sigma_1^{-1}$ decreases as the cube of σ_1^{-1}, and the relative error $(|r_{11}^{(1)}|^{-1} - \sigma_1^{-1})/\sigma_1^{-1}$ decreases as the square of $\sigma_2/\sigma_1 = 1/\sigma_1 = \sigma_1^{-1}$.

3.2 The Interior Singular Values

There are generalizations of the above bounds in terms of the norms of R_{11}^{-1} and R_{22}. That is, if R_{11} has dimension k and hence R_{22} dimension $n-k$, we can bound $\|(R_{11}^{(1)})^{-1}\| - \sigma_k^{-1}$ and $\|R_{22}^{(1)}\| - \sigma_{k+1}$.

For general k, the bounds provided for "Hybrid-III" by Chandrasekaran and Ipsen [2] are as follows:

$$\|R_{22}^{(0)}| \leq \sqrt{(k+1)(n-k)}\,\sigma_{k+1}, \quad (19)$$

$$\inf(R_{11}^{(0)}) \geq \frac{\sigma_k}{\sqrt{k(n-k+1)}}. \quad (20)$$

The general result is given in the following theorem.

Theorem 3.3 *Let A be an m-by-n matrix and let $\sigma_k(A) > \sigma_{k+1}(A)$. Let $R^{(0)}$ be the R-factor in the pivoted QR factorization of A, $A\Pi = Q^{(0)}R^{(0)}$ and let $R^{(1)}$ be the R-factor in the unpivoted QR factorization of $(R^{(0)})^T$, $(R^{(0)})^T = Q^{(1)}R^{(1)}$. Assume that the bounds (19) and (20) hold and that $\rho^{(0)} = \|R_{22}^{(0)}\|/inf(R_{11}^{(0)}) < 1$.*

Then for $j = 1, \ldots, n-k$,

$$\frac{\sigma_j(R_{22}^{(1)}) - \sigma_{k+j}}{\sigma_{k+j}} \leq \left(\frac{\sigma_{k+1}}{\sigma_k}\right)^2 \mathcal{O}\left(\frac{n^{\frac{5}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right), \quad (21)$$

and for $j = 1, \ldots, k$,

$$\frac{\sigma_j(R_{11}^{(1)})^{-1} - \sigma_j^{-1}}{\sigma_j^{-1}} \leq \left(\frac{\sigma_{k+1}}{\sigma_k}\right)^2 \mathcal{O}\left(\frac{n^{\frac{5}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right). \quad (22)$$

Proof: The proofs of inequalities (21) and (22) are similar to the proofs of Theorems 3.1 and 3.2, respectively. For (21), for example, Mathias and Stewart give us the first line

$$\sigma_j(R_{22}^{(1)}) - \sigma_{k+j} \leq \sigma_j(R_{22}^{(1)})\mathcal{O}\left(\frac{\|R_{12}^{(1)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2}\right),$$

and the rest follows as in the proof of Theorem 3.1. $\square$

Theorem 3.3 tells us that when there is a gap between σ_k and σ_{k+1}, the relative error in each of the singular values is quadratic in the gap ratio σ_{k+1}/σ_k. For the singular values at the gap, the absolute error is cubic in the singular value, as the following corollary spells out.

Corollary 3.4 *Under the assumptions of Theorem 3.3, we have*

$$\|R_{22}^{(1)}\| - \sigma_{k+1} \le \frac{\sigma_{k+1}^3}{\sigma_k^2} \mathcal{O}\left(\frac{n^{\frac{5}{2}} \|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2} \right), \tag{23}$$

and

$$\|(R_{11}^{(1)})^{-1}\| - \sigma_k^{-1} \le \frac{\sigma_{k+1}^2}{\sigma_k^3} \mathcal{O}\left(\frac{n^{\frac{5}{2}} \|R_{12}^{(0)}\|^2}{[1-(\rho^{(1)})^2][\inf(R_{11}^{(1)})]^2} \right). \tag{24}$$

Note from the bounds in Theorem 3.3 that the corollary represents the worst case. That is, the bounds are better as we go away from the gap. Theorem 3.3 helps explain why the QLP decomposition does so well even when the ratios between neighboring singular values are close to one. As long as there is one substantial gap somewhere in the singular values, convergence will be fast for all of them.

So we see how the pivoting is so important to the QLP decomposition. The pivoting in the first step roughly orders the embryonic singular values. In the theorems, this allows us to make the reasonable assumption (3) for $i = 0$. This ordering then allows the quadractic convergence factor in the second step.

In addition to providing this ordering, the pivoting in the first step also gives us a linear convergence factor (in the form of the RRQR bounds). Thus the rank-revealing nature of the pivoting furnishes one order of convergence, while the ordering nature of the pivoting sets the table for the second step to provide two additional orders of convergence.

Note that neither the assumption (3) nor the assumed RRQR bounds necessarily hold for the column pivoted QR decomposition. But the RRQR bounds often do, and (3) is more probable the larger the gap, so this analysis helps explain why QLP apparently does so well in practice. Without pivoting, the two-step algorithm not only fails to provide the linear convergence factor in the first step; more importantly, it often fails to provide the ordering needed for the quadratic convergence factor in the second step. Total failure is often seen with graded matrices. Pivoting is key.

We now illustrate Theorem 3.3 and Corollary 3.4 with a couple of examples. In each example we use a 100-by-100 matrix, let there be gap between σ_{50} and σ_{51}, and once again let the gap ratio σ_{51}/σ_{50} decrease from 10^{-1} initially to 10^{-5} in the fifth run.

The first example illustrates the bound (23). We fix the fifty largest singular values to be equally spaced from 1 to 10. The lowest fifty singular values are also equally spaced, but their range varies from 10^{-1} to 10^{-2} on the first run, from 10^{-2} to 10^{-3} on the second run, and so forth. As the gap ratio σ_{51}/σ_{50} thus decreases by a factor of 10 each run, we expect to see the absolute error $\|R_{22}^{(1)}\| - \sigma_{51}$ and the relative error $(\|R_{22}^{(1)}\| - \sigma_{51})/\sigma_{51}$ decrease by factors of 10^3 and 10^2, respectively. This is shown in Figure 4.

In the second example, we illustrate the bound (22). This time we fix the smallest fifty singular values to be equally spaced from 0.1 to 1. The largest fifty singular values range from 10 to 10^2, then 10^2 to 10^3, etc. As the gap ratio thus decreases by a factor of 10 each run, we expect to see the relative error in each of the first fifty singular values decrease by a factor of 10^2. This is borne out by Figure 5, in which we plot the relative error in σ_1^{-1} and σ_{40}^{-1}.

4 Convergence of the QLP Iteration

In section 2.2 we discussed taking the QLP decomposition beyond just the first two steps, resulting in an iteration (the QLP iteration) that converges to the singular values. In this section, we take two approaches to studying

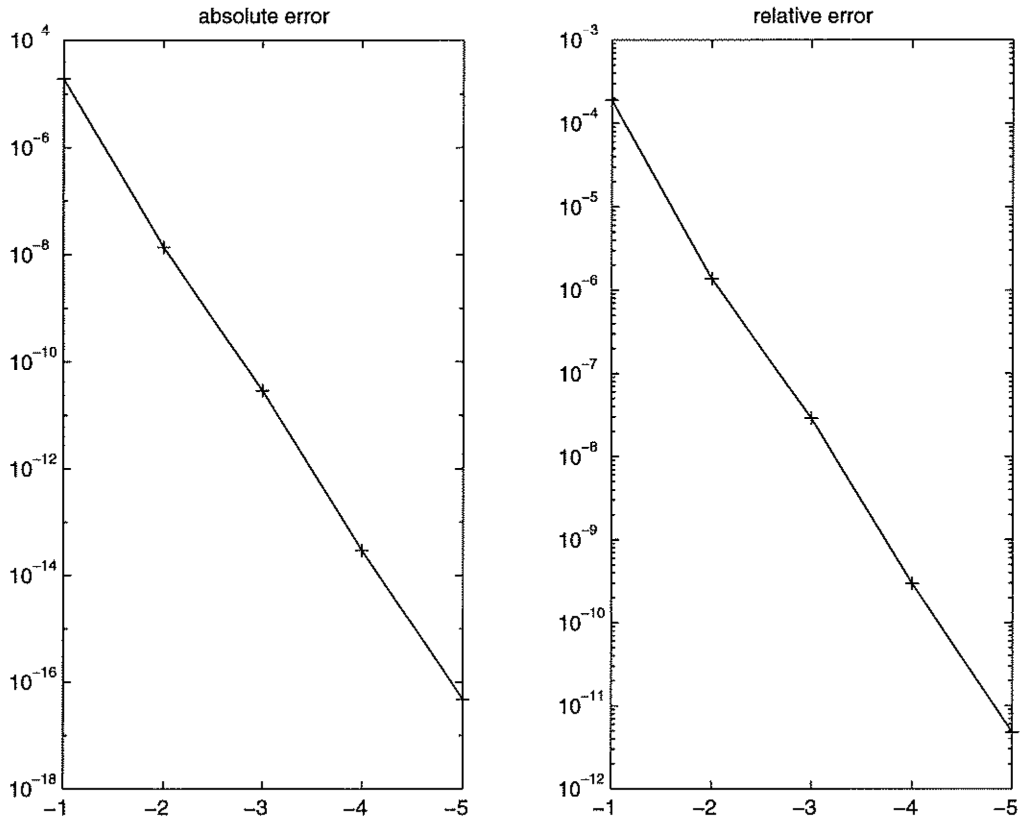

Figure 4: As σ_{51} decreases from 10^{-1} down to 10^{-5}, the absolute error $\|R_{22}^{(1)}\| - \sigma_{51}$ decreases as the cube of σ_{51}, and the relative error $(\|(R_{22}^{(1)}\| - \sigma_{51})/\sigma_{51}$ decreases as the square of $\sigma_{51}/\sigma_{50} = \sigma_{51}/1 = \sigma_{51}$.

the convergence of the QLP iteration.

4.1 More Iterations

In sections 3.1 and 3.2 we saw that with a gap between σ_k and σ_{k+1}, the QLP decomposition approximates all of the singular values with a relative error depending on the square of the gap ratio σ_{k+1}/σ_k. If we iterate beyond the first two steps (assuming that pivoting is used on the first step and no pivoting on each subsequent step), the error bounds improve by a quadratic

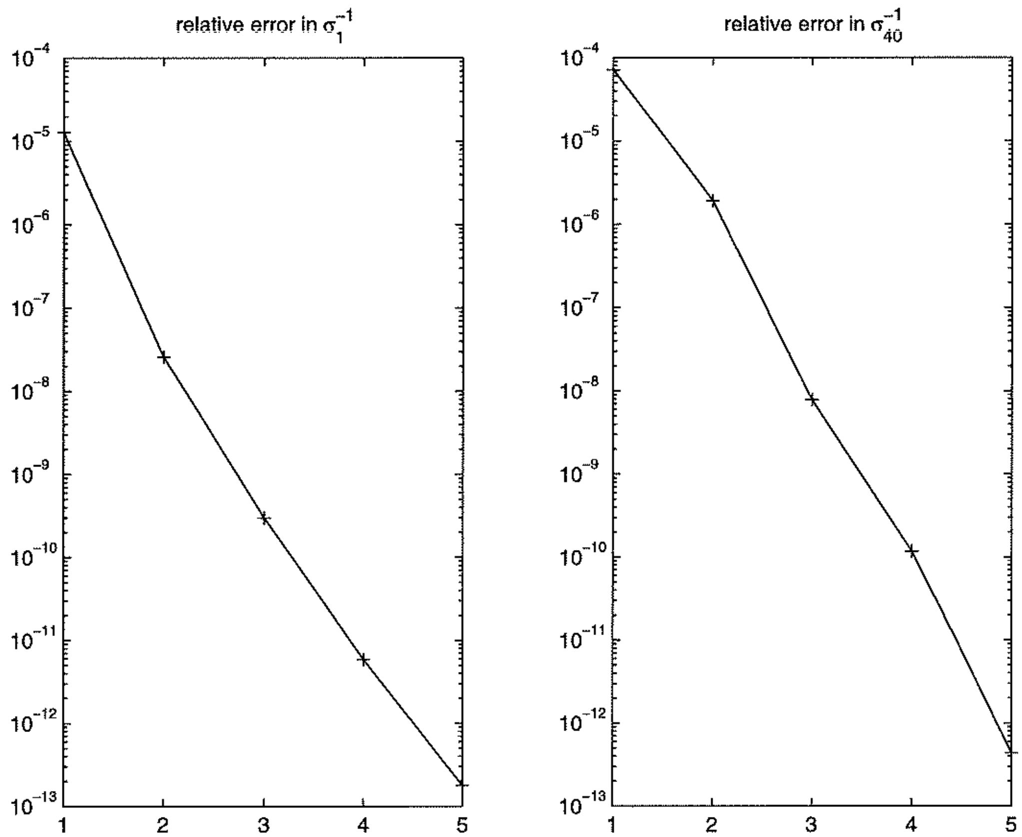

Figure 5: As σ_{50} increases from 10 up to 10^5, the relative error $(\sigma_1(R_{11}^{(1)})^{-1} - \sigma_1^{-1})/\sigma_1^{-1}$, plotted on the left, and the relative error $(\sigma_{40}(R_{11}^{(1)})^{-1} - \sigma_{40}^{-1})/\sigma_{40}^{-1}$, plotted on the right, both decrease as the square of of $\sigma_{51}/\sigma_{50} = 1/\sigma_{50}$.

factor in each step. To see this, we can simply apply the bound (4) at each iteration.

Theorem 4.1 *Let A be an m-by-n matrix and let $\sigma_k(A) > \sigma_{k+1}(A)$. Let $R^{(0)}$ be the R-factor in the pivoted QR factorization of A, $A\Pi = Q^{(0)}R^{(0)}$ and let $R^{(i)}$, $i \geq 1$, be the R-factor in the unpivoted QR factorization of $(R^{(i-1)})^T$, $(R^{(i-1)})^T = Q^{(i)}R^{(i)}$. Assume that the bounds (19) and (20) hold and that $\rho^{(0)} = \|R_{22}^{(0)}\|/inf(R_{11}^{(0)}) < 1$.*

Then for $j = 1, \ldots, n-k$,

$$\frac{\sigma_j(R_{22}^{(i)}) - \sigma_{k+j}}{\sigma_{k+j}} \leq \left(\frac{\sigma_{k+1}}{\sigma_k}\right)^{2i} \mathcal{O}\left(\frac{n^{\frac{4i+1}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(i)})^2][\inf(R_{11}^{(i)})]^2}\right), \tag{25}$$

and for $j = 1, \ldots, k$,

$$\frac{\sigma_j(R_{11}^{(i)})^{-1} - \sigma_j^{-1}}{\sigma_j^{-1}} \leq \left(\frac{\sigma_{k+1}}{\sigma_k}\right)^{2i} \mathcal{O}\left(\frac{n^{\frac{4i+1}{2}}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(i)})^2][\inf(R_{11}^{(i)})]^2}\right). \tag{26}$$

Proof:

The proofs are similar to those given before. For example, in proving the bound (25), the only real difference from the proof of Theorem 3.1 is inequality (9), different here because we are iterating. So the first two lines of the series of inequalities are as follows:

$$\begin{aligned} \sigma_j(R_{22}^{(i)}) - \sigma_{k+j} &\leq \sigma_j(R_{22}^{(i)})\mathcal{O}\left(\frac{\|R_{12}^{(i)}\|^2}{[1-(\rho^{(i)})^2][\inf(R_{11}^{(i)})]^2}\right) \\ &\leq \sigma_j(R_{22}^{(i)})\mathcal{O}\left(\frac{(\rho^{(0)})^{2i}\|R_{12}^{(0)}\|^2}{[1-(\rho^{(i)})^2][\inf(R_{11}^{(i)})]^2}\right). \end{aligned}$$

The rest of the proof proceeds as expected. □

4.2 Assymptotic Convergence of Individual Elements

We would now like to make some observations and conjectures on the rate of convergence of individual elements in the $R^{(i)}$ matrices as well as individual elements in the $U^{(i)}$ and $V^{(i)}$ matrices, where $R^{(i)} = U^{(i)}\Sigma(V^{(i)})^T$. Unfortunately, we cannot provide hard and fast proofs for everything, but we will describe what we see. It appears that the convergence of the diagonal elements of the $R^{(i)}$ depends on the convergence of the off-diagonal elements.

We will prove a statement concerning this dependence. The convergence of the off-diagonal elements appears complex, and we will give some observations.

We already know that the $R^{(i)}$ converge to a diagonal matrix whose elements are the singular values of A in decreasing order. So in discussing asymptotic rates of convergence of elements in $R^{(i)}$, we may assume that the diagonal elements are close to the σ_j and that all of the off-diagonal elements of $R^{(i)}$ are very small. With pivoting in the first step, convergence often reaches the asymptotic range quickly, after the first two or three iterations. Note that whether pivoting is used (in the first step only or on every step) or not, there will come an iteration after which pivoting ceases anyway. Hence, it would seem that asymptotic rates of convergence would not depend on whether we pivot. This intuition is wrong, however, for the acceleration in convergence afforded by pivoting extends in some cases even to the asymptotic convergence rates. We will comment on this.

Either Householder transformations or Givens rotations can be used to triangularize each $(R^{(i)})^T$. We will assume the use of Givens rotations, since their action is easier to track and visualize, and because they illuminate the rates of convergence for specific elements.

Let us look at how the elements of $(R^{(i)})^T$ could change while the matrix is being upper-triangularized via Givens rotations. The elements $(r^{(i)})^T_{21}$ through $(r^{(i)})^T_{n1}$ are first zeroed out, then the elements $(r^{(i)})^T_{32}$ through $(r^{(i)})^T_{n2}$ are zeroed out, and so on. We will assume that elements in each column are zeroed from top to bottom. Doing so gives us a useful analytical result. It is easy to verify that zeroing out the elements in this order leaves the submatrix consisting of the bottom $n-1$ rows and $n-1$ columns in lower-triangular form. This means that in zeroing out the second and subsequent columns, the mechanics look exactly the same.

When an element $(r^{(i)})^T_{jk}$ is zeroed out, the Givens rotation will change entries only in rows j and k. Take a look at Figure 6.

In Figure 6 we have assigned different letters to the nonzero elements that

$$
\begin{pmatrix}
x & x & x & x & x & x \\
 & \mathbf{x} & \mathbf{x} & \mathbf{x} & \mathbf{0} & \\
 & & x & & & \\
 & & x & x & & \\
 & \mathbf{x} & \mathbf{x} & \mathbf{x} & \mathbf{x} & \\
 & x & x & x & x & x
\end{pmatrix}
\qquad
\begin{pmatrix}
x & x & x & x & x & x \\
 & \mathbf{a} & \mathbf{c_1} & \mathbf{c_2} & \mathbf{0} & \\
 & & x & & & \\
 & & x & x & & \\
 & \mathbf{h} & \mathbf{d_1} & \mathbf{d_2} & \mathbf{b} & \\
 & x & x & x & x & x
\end{pmatrix}
$$

Figure 6: The element $(r^{(i)})^T_{52}$ is about to be zeroed out by a Givens rotation. On the left the elements which at this point are nonzero are marked with an x. The bold elements will be changed by the Givens rotation. On the right, depending on what type of change the Givens rotation will effect, the changing elements are assigned various letters.

change to indicate for each element what type of change the Givens rotation will effect. For example, the elements c_1 and c_2 will be changed in the same way. Figure 7 shows exactly how each element is changed.

$$
\begin{pmatrix}
a & \cdots & c_s & \cdots & 0 \\
h & \cdots & d_s & \cdots & b
\end{pmatrix}
\longrightarrow
\begin{pmatrix}
\sqrt{a^2+h^2} & \cdots & \frac{ac_s+hd_s}{\sqrt{a^2+h^2}} & \cdots & \frac{hb}{\sqrt{a^2+h^2}} \\
0 & \cdots & \frac{ad_s-hc_s}{\sqrt{a^2+h^2}} & \cdots & \frac{ab}{\sqrt{a^2+h^2}}
\end{pmatrix}
$$

Figure 7: A Givens rotation affects elements in two rows. Those elements are shown on the left, with letters indicating the type of change they will undergo. On the right, the elements appear as they are after the Givens rotation.

First of all, note that with every Givens rotation that is applied, the diagonal elements will always play the role of either a or b. In fact, the diagonal element $(r^{(i)})^T_{jj}$ will first play the role of b precisely $j-1$ times, namely when the $j-1$ elements to the left of it in row j are being zeroed. It will then play the role of a precisely $n-j$ times, namely when the $n-j$ elements below it in column j are being zeroed.

So the convergence of a diagonal element is completely determined by these $n-1$ Givens rotations. Looking at Figure 7, we see that whether the element is acting as a or as b for a given rotation, it receives an $\mathcal{O}(h^2)$

perturbation, where h is the off-diagonal element being zeroed by the given rotation. As all of these $n-1$ off-diagonal elements are converging to zero, it is obvious that asymptotically, only the largest of the $n-1$ elements will contribute significantly to the convergence of the diagonal element. The rest of the $n-1$ elements will only slightly alter the value of the diagonal element at each iteration. Since asymptotically, the largest of the $n-1$ elements will be the one(s) with the slowest asymptotic convergence rate, it is acceptable to ignore the contributions from all the other off-diagonal elements.

Let us make this more precise in the context of an example. We will write things in terms of a small quantity $\epsilon \ll \sigma_n$, and then we are free to discard lower order terms. Say we are looking at the essential limit $r_{11}^{(\infty)}$ of $r_{11}^{(i)}$, which will play the role of a for each of the $n-1$ Givens rotations that affect it each iteration. If $r_{k1}^{(i)}$ is asymptotically converging the slowest of all the elements in the first column below $r_{11}^{(i)}$, then we ignore the contribution of all the others. (We will later consider the case when two or more of these elements share the slowest convergence rate.) Let us assume that $|r_{k1}^{(i)}| < \epsilon$ for all $i \geq i_0$, some i_0 and all the other $|r_{j1}^{(i)}|, j \neq 1$ are less than ϵ^2 for all $i \geq i_0$. Let $\lambda_1, \ldots, \lambda_{n-1}$ be the asymptotic convergence rates of each of the $r_{j1}^{(i)}, 1 < j \leq n$. Then for $i \geq i_0$ we assume that $r_{j1}^{(i)}$ is converging with a rate of $\lambda_j + \delta_i$, where $|\delta_i| < \epsilon$. We will require i_0 sufficiently large so that the largest (slowest) rate satisfies $\lambda_k + \epsilon < 1$. Letting $h = r_{k1}^{(i)}$, we look back at Figure 7 and see that for $i \geq i_0$,

$$\begin{aligned} r_{11}^{(i)} &= \sqrt{(r_{11}^{(i-1)})^2 + \sum_{j=2}^{n} (r_{j1}^{(i)})^2} \\ &= \sqrt{(r_{11}^{(i-1)})^2 + h^2} + \mathcal{O}(\epsilon^4), \end{aligned}$$

We see that of all the $n-1$ off-diagonal terms that could contribute to $r_{11}^{(i)}$, only $h = r_{k1}^{(i)}$ makes an $\mathcal{O}(\epsilon^2)$ contribution. All the rest effect only $\mathcal{O}(\epsilon^4)$ perturbations. So in a similar fashion we have

$$
\begin{aligned}
r_{11}^{(i+1)} &= \sqrt{(r_{11}^{(i-1)})^2 + h^2 + [(\lambda_k + \delta_{i+1})h]^2 + \mathcal{O}(\epsilon^4)} \\
r_{11}^{(i+2)} &= \sqrt{(r_{11}^{(i-1)})^2 + h^2 + [(\lambda_k + \delta_{i+1})h]^2 + [(\lambda_k + \delta_{i+1})(\lambda_k + \delta_{i+2})h]^2 + \mathcal{O}(\epsilon^4)} \\
&\cdots \\
r_{11}^{(\infty)} &= \sqrt{(r_{11}^{(i-1)})^2 + h^2 + \sum_{j=0}^{\infty} \left[\left(\prod_{a=i_0+1}^{i_0+1+j} (\lambda_k + \delta_a) \right) h \right]^2 + \mathcal{O}(\epsilon^4)}. \qquad (27)
\end{aligned}
$$

In equation (27), there might be a question as to whether we can actually sweep the contributions of all of the $n-2$ other off-diagonal elements into the $\mathcal{O}(\epsilon^4)$, since each of these elements contributes an infinite series similar to that of $h = r_{k1}^{(i)}$. The contribution from a generic one $\hat{h} = r_{b1}^{(i)}$ can be bounded as follows:

$$
\begin{aligned}
\sum_{j=0}^{\infty} \left[\left(\prod_{a=i_0+1}^{i_0+1+j} (\lambda_b + \delta_a) \right) \hat{h} \right]^2 &\leq \sum_{j=0}^{\infty} \left[(\lambda_b \pm \epsilon)^{j+1} \hat{h} \right]^2 \\
&= \hat{h}^2 (\lambda_b + \epsilon)^2 \sum_{j=0}^{\infty} (\lambda_b \pm \epsilon)^{2j} \\
&= \hat{h}^2 (\lambda_b \pm \epsilon)^2 \frac{1}{1 - (\lambda_b \pm \epsilon)^2} \\
&= \hat{h}^2 (\lambda_b \pm \epsilon)^2 (1 + \mathcal{O}\left((\lambda_b \pm \epsilon)^2\right) \\
&= \mathcal{O}(\epsilon^4),
\end{aligned}
$$

since $\hat{h} = \mathcal{O}(\epsilon^4)$ and $\lambda_b + \epsilon < \lambda_k + \epsilon < 1$ by assumption when $i \geq i_0$. So equation (27) holds.

Now we have a feel for the proof of the following theorem, which concerns the asymptotic convergence rates of the diagonal elements.

Theorem 4.2 *Let $\epsilon \ll \sigma_n$. Let i_0 be an integer sufficiently large so that for $i \geq i_0$ the following hold:*

1. *For each $1 \le s \le n$, denote by h_s the element having the slowest (i.e., largest) asymptotic convergence rate of all the elements in the same row or column as $r_{ss}^{(i)}$, namely, $(r^{(i)})^T_{s,1}, \ldots, (r^{(i)})^T_{s,s-1}$ and $(r^{(i)})^T_{s+1,s}, \ldots, (r^{(i)})^T_{n,s}$, and denote the asymptotic convergence rate of h_s by λ_s. If more than one of these $n-1$ elements share the same slowest asymptotic convergence rate, denote them by h_{s_1}, h_{s_2}, etc.*

2. *Assume the rate of convergence of each of these $n-1$ elements is within ϵ of its asymptotic convergence rate. So for example, the rate of the element h_s at iteration $i \ge i_0$ is $\lambda_s + \delta_i$, where $|\delta_i| < \epsilon$. Similarly for the rest of these $n-2$ elements.*

3. *Assume all of the off-diagonal elements of the matrix $R^{(i)}$ are less than ϵ^2 in absolute value, except for the $h_{s_j}, 1 \le s \le n$, which need only be less than ϵ in absolute value.*

4. *Assume that each of the diagonal elements $r_{ss}^{(i)}, 1 \le s \le n$ is within ϵ^2 of σ_s.*

5. *Assume $\lambda_s + \epsilon < 1$, for $1 \le s \le n$.*

Then the convergence rate of $r_{ss}^{(i)}$ is $\lambda_s^2 + \mathcal{O}(\epsilon)$.

Proof: Without loss of generality, we assume that the diagonal elements of the $R^{(i)}$ are positive.

First let $s = 1$.

Say there is only one element h_1 among the $n-1$ off-diagonal elements in the first column of $R^T_{(i)}$ that converges with the slowest asymptotic rate λ_1. We have already derived formulae for $r_{11}^{(i)}$ and $r_{11}^{(\infty)}$ above. Let $P = \prod_{a=i_0+1}^{i_0+1+j}(\lambda_1 + \delta_a)$. We have

$$
\begin{aligned}
r_{11}^{(\infty)} - r_{11}^{(i)} &= \sqrt{(r_{11}^{(i-1)})^2 + h_1^2 + \sum_{j=0}^{\infty} (Ph_1)^2} - \sqrt{(r_{11}^{(i-1)})^2 + h_1^2} + \mathcal{O}(\epsilon^4) \\
&= |r_{11}^{(i-1)}| \left(1 + \frac{1}{2} \frac{h_1^2 + \sum_{j=0}^{\infty} (Ph_1)^2}{(r_{11}^{(i-1)})^2} - 1 - \frac{1}{2} \frac{h_1^2}{(r_{11}^{(i-1)})^2} \right) + \mathcal{O}(\epsilon^4) \\
&= \frac{1}{2} \frac{\sum_{j=0}^{\infty} (Ph_1)^2}{r_{11}^{(i-1)}} + \mathcal{O}(\epsilon^4).
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
r_{11}^{(\infty)} - r_{11}^{(i+1)} &= \sqrt{(r_{11}^{(i-1)})^2 + h_1^2 + \sum_{j=0}^{\infty} (Ph_1)^2} \\
&\qquad - \sqrt{(r_{11}^{(i-1)})^2 + h_1^2 + ((\lambda_1 + \delta_{i_0+1})h_1)^2} + \mathcal{O}(\epsilon^4) \\
&= \frac{1}{2} \frac{\sum_{j=1}^{\infty} (Ph_1)^2}{r_{11}^{(i-1)}} + \mathcal{O}(\epsilon^4).
\end{aligned}
$$

So the asymptotic rate of convergence of $r_{11}^{(i)}$ is

$$
\begin{aligned}
\frac{r_{11}^{(\infty)} - r_{11}^{(i+1)}}{r_{11}^{(\infty)} - r_{11}^{(i)}} &= \frac{\frac{1}{2} \frac{\sum_{j=1}^{\infty} (Ph_1)^2}{r_{11}^{(i-1)}}}{\frac{1}{2} \frac{\sum_{j=0}^{\infty} (Ph_1)^2}{r_{11}^{(i-1)}}} + \mathcal{O}(\epsilon^4) \\
&= \frac{\sum_{j=1}^{\infty} (Ph_1)^2}{\sum_{j=0}^{\infty} (Ph_1)^2} + \mathcal{O}(\epsilon^4).
\end{aligned}
$$

To bound this, we can set all of the δ_a in $P = \prod_{a=i_0+1}^{i_0+1+j} (\lambda_1 + \delta_a)$ to their maximum values, obtaining

$$
\begin{aligned}
\frac{h_1^2\sum_{j=1}^{\infty}(\lambda_1 \pm \epsilon)^{2j+4}}{h_1^2\sum_{j=0}^{\infty}(\lambda_1 \pm \epsilon)^{2j+2}} + \mathcal{O}(\epsilon^4) &= \frac{(\lambda_1 \pm \epsilon)^2\frac{1}{1-(\lambda_1\pm\epsilon)^2}}{\frac{1}{1-(\lambda_1\pm\epsilon)^2}} + \mathcal{O}(\epsilon^4) \\
&= (\lambda_1 \pm \epsilon)^2 + \mathcal{O}(\epsilon^4) \\
&= \lambda_1^2 + \mathcal{O}(\epsilon)
\end{aligned}
$$

Now consider the case when $s = n$. The only Givens rotations that affect $(r^{(i)})^T_{nn}$ are the ones that zero $(r^{(i)})^T_{n1}, \ldots, (r^{(i)})^T_{n,n-1}$. Again assume there is precisely one of these elements, say $(r^{(i)})^T_{k1}$, which has the slowest asymptotic convergence rate. Set $h_n = (r^{(i)})^T_{k1}$, and let its asymptotic convergence rate be λ_n.

We can look at Figure 7 to see how things will work. The b in Figure 7 will be here $r_{nn}^{(i-1)}$, and at each iteration this will be multiplied by $a/\sqrt{a^2+h^2}$. The element a will be $r_{kk}^{(i)}$ and the element h will be h_n, both of which change at each iteration. All the elements in the last row of $(R^{(i)})^T$ other than h_n will cause only $\mathcal{O}(\epsilon^4)$ perturbations as before. This time we let $P = \prod_{a=i_0+1}^{i_0+j}(\lambda_n + \delta_a)$, and we use the convention that $\prod_{a=i_0+1}^{i_0}(\lambda_n + \delta_a) = 1$. We have

$$
\begin{aligned}
r_{nn}^{(\infty)} - r_{nn}^{(i)} &= r_{nn}^{(i-1)} \prod_{j=0}^{\infty} \frac{r_{kk}^{(i+j)}}{\sqrt{(r_{kk}^{(i+j)})^2 + (Ph_n)^2}} \\
&\qquad\qquad - r_{nn}^{(i-1)} \frac{r_{kk}^{(i)}}{\sqrt{(r_{kk}^{(i)})^2 + h_n^2}} + \mathcal{O}(\epsilon^4) \qquad (28) \\
&= r_{nn}^{(i-1)} \prod_{j=0}^{\infty} \left[1 - \frac{1}{2} \frac{h_n^2}{(r_{kk}^{(i+j)})^2} P^2 \right] \\
&\qquad\qquad - r_{nn}^{(i-1)} \left(1 - \frac{1}{2} \frac{h_n^2}{(r_{kk}^{(i)})^2} \right) + \mathcal{O}(\epsilon^4) \\
&= r_{nn}^{(i-1)} \left(1 - \frac{h_n^2}{2} \sum_{j=0}^{\infty} \frac{P^2}{(r_{kk}^{(i+j)})^2} - 1 + \frac{1}{2} \frac{h_n^2}{(r_{kk}^{(i)})^2} \right) + \mathcal{O}(\epsilon^4) \\
&= -r_{nn}^{(i-1)} \frac{h_n^2}{2} \sum_{j=1}^{\infty} \frac{P^2}{(r_{kk}^{(i+j)})^2} + \mathcal{O}(\epsilon^4).
\end{aligned}
$$

This leads to

$$
\frac{r_{nn}^{(\infty)} - r_{nn}^{(i+1)}}{r_{nn}^{(\infty)} - r_{nn}^{(i)}} = \frac{-r_{nn}^{(i-1)} \frac{h_n^2}{2} \sum_{j=2}^{\infty} \frac{P^2}{(r_{kk}^{(i+j)})^2}}{-r_{nn}^{(i-1)} \frac{h_n^2}{2} \sum_{j=1}^{\infty} \frac{P^2}{(r_{kk}^{(i+j)})^2}} + \mathcal{O}(\epsilon^4). \qquad (29)
$$

To bound this, we once again set all of the δ_a in $P = \prod_{a=i_0+1}^{i_0+j} (\lambda_n + \delta_a)$ to their maximum values. We also make use of the fact that the diagonal elements are within ϵ^2 of the singular values to which they are converging. This gives us

$$
\frac{\sum_{j=2}^{\infty} \frac{(\lambda_n \pm \epsilon)^{2j+2}}{\sigma_k \pm \epsilon^2}}{\sum_{j=1}^{\infty} \frac{(\lambda_n \pm \epsilon)^{2j}}{\sigma_k \pm \epsilon^2}} + \mathcal{O}(\epsilon^4) = \lambda_n^2 + \mathcal{O}(\epsilon). \qquad (30)
$$

Finally let $s = l$. The only Givens rotations that affect $(r^{(i)})^T_{ll}$ are the ones that zero $(r^{(i)})^T_{l1}, \dots, (r^{(i)})^T_{l,l-1}$, and the ones that zero $(r^{(i)})^T_{l+1,l}, \dots, (r^{(i)})^T_{n,l}$. Assume one of these has the slowest asymptotic convergence rate. Then the proof for either the $s = 1$ case or the $s = n$ case carries over.

If in any of the cases more than one of these elements converge at λ_s, the same proofs hold with minor modifications. For example, say that both $(r^{(i)})^T_{pl}$ and $(r^{(i)})^T_{lq}$ converge at the rate λ_s. Denote $r^{(i-1)}_{ql}$ by h_{s_1} and $r^{(i-1)}_{lp}$ by h_{s_2}. So looking once again at Figure 7, $(r^{(i)})^T_{ll}$ plays the role of b with $h = h_{s_1}$ and later plays the role of a with $h = h_{s_2}$. Hence $r^{(\infty)}_{ll}$, the analoge of $r^{(\infty)}_{nn}$ in equation (28), is

$$\begin{aligned} & r^{(i-1)}_{ll} \prod_{j=0}^{\infty} \sqrt{\left(\frac{r^{(i+j)}_{qq}}{\sqrt{(r^{(i+j)}_{qq})^2 + P_1^2}} \right)^2 + \left(\frac{P_2}{r^{(i-1)}_{ll}} \right)^2} + \mathcal{O}(\epsilon^4) \\ & \quad = r^{(i-1)}_{ll} \prod_{j=0}^{\infty} \sqrt{1 - \left(\frac{P_1}{r^{(i+j)}_{qq}} \right)^2 + \left(\frac{P_2}{r^{(i-1)}_{ll}} \right)^2} + \mathcal{O}(\epsilon^4) \\ & \quad = r^{(i-1)}_{ll} \left(1 + \frac{1}{2} \sum_{j=0}^{\infty} \left[\frac{P_2^2}{(r^{(i-1)}_{ll})^2} - \frac{P_1^2}{(r^{(i+j)}_{qq})^2} \right] \right) + \mathcal{O}(\epsilon^4). \end{aligned}$$

Here we used $P_1 = \left(\prod_{a=i_0+1}^{i_0+j} (\lambda_l + \delta_a) \right) h_{l_1}$ and $P_2 = \left(\prod_{a=i_0+1}^{i_0+j} (\lambda_l + \delta_a) \right) h_{l_2}$. The proof then continues as in the other cases. (At the step where we bound, as in going from (29) to (30), letting each $\delta_a = \pm\epsilon$ and $r^{(i+j)}_{qq} = \sigma_q \pm \epsilon^2$, everything inside the sum in the numerator except for $(\lambda_l \pm \epsilon)^{2j+2}$ factors out and cancels with the identical terms in the denominator.) For other cases in which multiple elements share the slowest asymptotic convergence rate, the proofs are similar. □

Theorem 4.2 basically says that the asymptotic rates of convergence of the diagonal elements are the squares of the rates of the off-diagonal elements. More precisely, the rate of a diagonal element $r^{(i)}_{ss}$ is equal to the square of the slowest rate among all the off-diagonal elements in the union of row s and column s. So we now know the asymptotic convergence rates of the diagonal elements given the asymptotic convergence rates of the off-diagonal elements, to which we now turn.

We can use Mathias and Stewart's results again here. Recall that under the assumption (3) we have

$$\|R_{12}^{(i+1)}\| \leq \rho^{(i)} \|R_{12}^{(i)}\|.$$

For sufficiently large i, if $\sigma_k \neq \sigma_{k+1}$, then (3) holds. So we know that every element in the off-diagonal block $R_{12}^{(i)}$ must converge asymptotically with a rate $\rho \leq \sigma_{k+1}/\sigma_k$. So we have the following bound on the convergence of off-diagonal elements.

Lemma 4.3 *For $s < t$ let λ_{st} be the asymptotic rate of convergence of $r_{st}^{(i)}$. If $\sigma_s \neq \sigma_t$ then $\lambda_{st} \leq \sigma_t/\sigma_s$.*

Combining Lemma 4.3 and Theorem 4.2 gives us the following bound.

Theorem 4.4 *Let λ_{ss} be the asymptotic convergence rate of $r_{ss}^{(i)}$. If $\sigma_{s-1} > \sigma_s > \sigma_{s+1}$, then $\lambda_{ss} \leq \max((\sigma_s/\sigma_{s-1})^2, (\sigma_{s+1}/\sigma_s)^2)$.*

This bound is tight. In addition, there are many times when the convergence is much faster. Examples given later will illustrate both of these facts.

Having achieved a bound on the convergence of off-diagonal (and therefore diagonal) elements, we can make a few observations on precise asymptotic convergence rates.

We can easily see what happens when A is a 2-by-2 matrix with distinct singular values, for the passage from $(R^{(i)})^T$ to $R^{(i+1)}$ is illustrated by Figure 7 (just look at the first and last columns). Since we are in the asymptotic range, h is very small, and the diagonal elements a and b are close to σ_1 and

σ_2. They experience $\mathcal{O}(h^2)$ perturbations and thus remain close to σ_1 and σ_2. If we denote the off-diagonal element of $R^{(i+1)}$ by h', then the rate of convergence is $h'/h = b/\sqrt{a^2+h^2}$. This is an $\mathcal{O}(h^2)$ perturbation of $b/|a|$, so we see that the asymptotic rate of convergence is σ_2/σ_1. Note that if the singular values are equal, the convergence is painfully slow.

The elements in a general n-by-n matrix are experiencing two types of effects from the various Givens rotations, and these can be seen once again in Figure 7. Once each iteration, a given off-diagonal element x will be zeroed out. That is, it will play the role of h in Figure 7. Let us say that x here receives a *zeroing contribution* to its convergence. Typically several times each iteration, the element x will be affected by a Givens rotation though it is not being zeroed out. That is, it will play the role of c_s or d_s in Figure 7. Let us say that x is here receiving a *nonzeroing contribution.* (The only off-diagonal element never to receive a nonzeroing contribution to its convergence is $r_{1n}^{(i)}$. It plays the role of h precisely once, with $a = r_{11}^{(i)}$ and $b = r_{nn}^{(i)}$, and therefore its asymptotic rate of convergence is σ_n/σ_1, the reciprocal of the condition number.)

We already understand the effect of a zeroing contribution. Indeed, if only zeroing contributions were in play, then the situation in an n-by-n matrix would be a simple generalization of the 2-by-2 case: $r_{st}^{(i)}$ would asymptotically converge at a rate of σ_t/σ_s if $\sigma_s > \sigma_t$ and painfully slowly if the singular values were equal. We know this is not the case from empirical observation. First, repeated singular values tend to speed up convergence, and speed it up a great deal. See Figure 8. As we know that this is not coming from zeroing contributions, we have our first hint that nonzeroing contributions not only sometimes play a significant role but can be agents of very fast convergence. Second, nonzeroing contributions sometimes play a significant role (and bring about fast convergence) even in matrices with all distinct singular values.

To get a feel for how this fast convergence might happen, let us take a closer look at a nonzeroing contribution. Say an off-diagonal element in question plays the role of c_s in Figure 7 under a certain Givens rotation. If

$$\begin{pmatrix} \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{2} & \frac{1}{50} & \frac{1}{10} \\ & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{25} & \frac{1}{5} \\ & & \underset{\frac{1}{25}}{\mathbf{5}} & \frac{1}{5} \\ & & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}$$

Figure 8: This is one observed convergence pattern for a matrix having singular values 10, 5, 5, and 1, written in bold on the diagonal for reference. In place of each off-diagonal element is its asymptotic convergence rate. Beneath each diagonal element is its rate. (Note that the diagonal convergence is as Theorem 4.2 states.) We know that nonzeroing contributions must be significant simply because of the repeated singular values. But look at how they speed up the convergence. The rates of the elements $r_{13}^{(i)}$ and $r_{23}^{(i)}$ are faster than the ratio of any two singular values in the matrix. We would expect $r_{13}^{(i)}$ to converge at a rate of $5/10 = 1/2$ were zeroing contributions dominating. Note that its (much) faster rate here allows the diagonal element $r_{33}^{(i)}$ to have a faster rate as well, $(1/5)^2$ instead of $(1/2)^2$.

we denote its value after this rotation is applied by c_s', then the convergence rate is

$$\frac{c_s'}{c_s} = \frac{ac_s - hd_s}{c_s\sqrt{a^2+h^2}} \approx \pm 1 + \frac{hd_s}{|a|c_s}, \tag{31}$$

where the sign is the sign of a and is different from the sign of the fraction. (The signs of the elements is a topic we will not pursue here.) Keeping in mind that a is large, being on the diagonal, we see that the convergence rate will be about ± 1 if h, c_s, and d_s are all about the same size. That is, the Givens rotation will not significantly affect c_s. But what if in equation (31), the element c_s is much smaller than h and d_s? Then the ratio $hd_s/|a|c_s$ is no longer close to zero, and the contribution from this Givens rotation is significant.

So where there are disparities in the sizes of off-diagonal elements, nonzeroing contributions can come into play. To see how they can speed convergence, again consider the asymptotic convergence rate of c_s given by equation (31). For this rate to be constant, c_s must converge at the product of the rates of h and d_s. (The diagonal element a is large and constant for our purposes here.) This is what we seem to observe with, for example, repeated singular values. See Figure 9.

$$\begin{pmatrix} \sim 10 & x & x & x \\ & a \sim 5 & \mathbf{c} & \mathit{0} \\ & & \sim 5 & \\ & h & d & b \sim 1 \end{pmatrix} \qquad \begin{pmatrix} \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{2} & \frac{1}{50} & \frac{1}{10} \\ & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{25} & \frac{1}{5} \\ & & \underset{\frac{1}{25}}{\mathbf{5}} & \frac{1}{5} \\ & & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}$$

Figure 9: The example of Figure 8 revisited. In the matrix on the left, the last Givens rotation applied gave the element $r_{23}^{(i)}$ (now marked $\mathbf{c}$) almost no zeroing contribution to its convergence because of the repeated singular values 5 and 5. Now the element marked h is being zeroed, and $\mathbf{c}$ will receive a nonzeroing contribution, because in this example it is very small compared to d and h. From the discussion, $\mathbf{c}$ will converge at the product of the rates of h and d. In this example, h and d are being affected by primarily zeroing contributions, so they are each converging at the rate of $1/5$. So $\mathbf{c}$ should be converging at the rate of their product, $1/25$. The matrix on the right, which shows the rates again for reference, confirms this.

Before we state an observation concerning the asymptotic rate of convergence for off-diagonal elements, let us set some notation to handle repeated singular values. Say that $\sigma_k > \sigma_{k+1} = \cdots = \sigma_{k+p} > \sigma_{k+p+1}$. Set $\rho_{k+1}^{(u)} = \cdots = \rho_{k+p}^{(u)} \equiv \sigma_k/\sigma_{k+1}$, and set $\rho_{k+1}^{(l)} = \cdots = \rho_{k+p}^{(l)} \equiv \sigma_{k+p}/\sigma_{k+p+1}$.

We already know from Lemma 4.3 that if $s < t$, then the slowest the asymptotic rate of convergence of $r_{st}^{(i)}$ can be is σ_s/σ_t, assuming the singular values are not equal. For the general case, we here catalogue some of the possible asymptotic rates of convergence for the off-diagonal elements based on our empirical observations.

1. If neither σ_s nor σ_t is a repeated singular value, then $r_{st}^{(i)}$ often converges at the rate σ_s/σ_t, but can behave as if σ_s and/or σ_t were repeated, as described in the following.

2. If $\sigma_s = \sigma_t$ then $r_{st}^{(i)}$ often converges at either the rate $(\rho_s^{(u)})^2$ or the rate $(\rho_s^{(l)})^2$. Actually, a sequence of j equal singular values defines a j-by-j block, all of whose off-diagonal elements converge at the same rate.

3. If $\sigma_s \neq \sigma_t$ and one or both are repeated singular values, then $r_{st}^{(i)}$ often converges at the rate σ_s/σ_t, $(\rho_s^{(u)})^2\sigma_s/\sigma_t$, $(\rho_s^{(l)})^2\sigma_s/\sigma_t$, $(\rho_t^{(u)})^2\sigma_s/\sigma_t$, or $(\rho_t^{(l)})^2\sigma_s/\sigma_t$.

In Figure 10 are the asymptotic convergence rates for a few example matrices with given singular values, and (1) through (3) above are represented.

In the 3-by-3 matrix, off-diagonal elements receive only zeroing contributions, so the asymptotic convergence rate of $r_{st}^{(i)}$ is σ_t/σ_s. This is not so for the 6-by-6 matrix. Although it has all singular values distinct, the element $r_{45}^{(i)}$ is receiving nonzeroing contributions, converging at a rate of $1/50 \ll 1/2 = 5/10 = \sigma_5/\sigma_4$. The three 4-by-4 matrices have the same singular values, 10, 5, 5, and 1, but converge differently. From these three matrices, we see that asymptotic (and not just initial) convergence is dependent not only on the singular values but also on the entries in the original matrix. It is also dependent on whether pivoting is used or not. In the examples we looked at, pivoting on the first step of the QLP iteration tended to give faster asymptotic convergence rates than not pivoting. (For example, for 4-by-4 matrices having singular values 10, 5, 5, and 1, the rates on the right are far more likely to occur when pivoting is used on the first step. Recall that it is discrepancies in the sizes of off diagonal elements that bring into play nonzeroing contributions and therefore faster convergence. Pivoting seems to encourage this.)

From studying the asymptotic convergence rates, we now have two insights into how the convergence of individual elements in the pivoted QLP

$$
\begin{pmatrix} \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{2} & \frac{1}{10} \\ & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{5} \\ & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}
\quad
\begin{pmatrix} \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{2} & \frac{1}{2} & \frac{1}{10} \\ & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{4} & \frac{1}{5} \\ & & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{5} \\ & & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}
\quad
\begin{pmatrix} \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{2} & \frac{1}{2} & \frac{1}{10} \\ & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{4} & \frac{1}{20} \\ & & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{5} \\ & & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}
\quad
\begin{pmatrix} \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{2} & \frac{1}{50} & \frac{1}{10} \\ & \underset{\frac{1}{4}}{\mathbf{5}} & \frac{1}{25} & \frac{1}{5} \\ & & \underset{\frac{1}{25}}{\mathbf{5}} & \frac{1}{5} \\ & & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}
$$

$$
\begin{pmatrix} \underset{\frac{1}{9}}{\mathbf{300}} & \frac{1}{3} & \frac{1}{6} & \frac{1}{30} & \frac{1}{60} & \frac{1}{300} \\ & \underset{\frac{1}{4}}{\mathbf{100}} & \frac{1}{2} & \frac{1}{10} & \frac{1}{20} & \frac{1}{100} \\ & & \underset{\frac{1}{4}}{\mathbf{50}} & \frac{1}{5} & \frac{1}{10} & \frac{1}{50} \\ & & & \underset{\frac{1}{25}}{\mathbf{10}} & \frac{1}{50} & \frac{1}{10} \\ & & & & \underset{\frac{1}{25}}{\mathbf{5}} & \frac{1}{5} \\ & & & & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}
\quad
\begin{pmatrix} \underset{\frac{1}{4}}{\mathbf{20}} & \frac{1}{4} & \frac{1}{8} & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{20} \\ & \underset{\frac{1}{4}}{\mathbf{20}} & \frac{1}{2} & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{20} \\ & & \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{4} & \frac{1}{8} & \frac{1}{8} & \frac{1}{8} & \frac{1}{40} \\ & & & \underset{\frac{1}{4}}{\mathbf{10}} & \frac{1}{8} & \frac{1}{8} & \frac{1}{8} & \frac{1}{10} \\ & & & & \underset{\frac{1}{16}}{\mathbf{5}} & \frac{1}{16} & \frac{1}{16} & \frac{1}{5} \\ & & & & & \underset{\frac{1}{16}}{\mathbf{5}} & \frac{1}{16} & \frac{1}{5} \\ & & & & & & \underset{\frac{1}{16}}{\mathbf{5}} & \frac{1}{5} \\ & & & & & & & \underset{\frac{1}{25}}{\mathbf{1}} \end{pmatrix}
$$

Figure 10: Here are six sample matrices with their singular values given in bold on the diagonal. Off-diagonal elements are replaced by their asymptotic convergence rates in italics, and diagonal elements have their rates written below them. Concerning the off-diagonal rates, see the text for some discussion and also notice that each of (1) through (3) in the catalogue on page 37 is represented. Concerning the diagonal rates, notice that Theorem 4.2 is everywhere verified: the rate of $r_{ss}^{(i)}$ is equal to the square of the slowest (i.e., largest) rate among all the off-diagonal elements in row s and column s.

decomposition can be fast in the presence of repeated singular values. One, in the presence of repeated singular values, convergence tends to be faster than with distinct singular values. Significant nonzeroing contributions cause the faster convergence, and they are necessarily a part of the landscape when repeated singular values are around. They are optional (and indeed less common, at least in our observations) when singular values are distinct. Although these are only asymptotic results, the pivoting can often speed up the approach of the asymptotic range.

Two, recall that nonzeroing contributions are caused by large discrepancies in the sizes of off-diagonal elements and that pivoting tends to encourage this. The discrepancies are thus starting to form in the first step, and where there are discrepancies, there are some small off-diagonal elements already.

These observations apply not only when there are repeated singular values but also when the ratios between neighboring singular values are close to one. Indeed, prior to reaching the asymptotic range, the QLP iteration cannot distinguish identical singular values from those that are merely close. The slow convergence in the latter case is an asymptotic phenomenon. Early on the convergence is fast.

It is also interesting to note that when faster convergence occurs, the convergence of the matrices of singular vectors of the $R^{(i)}$ is also faster. For example, Chadrasekaran and Ipsen [3] show that the convergence to zero of the angle between the first k columns of $U^{(i)}$ (and also of $V^{(i)}$) and the first k columns of the identity matrix is bounded by σ_{k+1}/σ_k. We would like to point out that faster rates are attainable when nonzeroing contributions are at work. For example, in the 6-by-6 matrix in Figure 10, we noticed that the element in the (4,5) position is converging fast, at a rate of 1/50 instead of at the "expected" rate of $\sigma_5/\sigma_4 = 5/10 = 1/2$. Correspondingly, the first four columns of $U^{(i)}$ and of $V^{(i)}$ are converging not at a rate of 1/2 but at a rate of 1/10.

Finally, we note that there are patterns in the asymptotic convergence of individual off-diagonal elements of the $U^{(i)}$ and $V^{(i)}$ similar to the patterns seen in the $R^{(i)}$.

5 Conclusion

We have studied Stewart's pivoted QLP decomposition, which represents the first two steps in an algorithm which approximates the SVD. The matrix $A\Pi_0$ is first factored as $A\Pi_0 = QR$, and then the matrix $R^T\Pi_1$ is factored as $R^T\Pi_1 = PL^T$, resulting in $A = Q\Pi_1 LP^T\Pi_0^T$, with Q and P orthogonal, L lower-triangular, and Π_0 and Π_1 permutation matrices. Stewart noted that the diagonal elements of L approximate the singular values of A with surprising accuracy, and we have provided mathematical justification for this phenomenon.

Specifically, we showed that if there is a gap between σ_k and σ_{k+1}, partition the matrix L into diagonal blocks L_{11} and L_{22} and off-diagonal block L_{21}, where L_{11} is k-by-k. We show that the convergence of $(\sigma_j(L_{11})^{-1} - \sigma_j^{-1})/\sigma_j^{-1}$ for $j = 1, \ldots, k$, and of $(\sigma_j(L_{22}) - \sigma_{k+j})/\sigma_{k+j}$, for $j = 1, \ldots, n-k$ are all quadratic in the gap ratio σ_{k+1}/σ_k. Hence the pivoted QLP decomposition will probably approximate the singular values very well when there is at least one large gap anywhere in the singular values, even if most of the other ratios between neighboring singular values are close to one.

The "worst case" for the bounds are at the gap, where the absolute errors $\|L_{11}^{-1}\| - \sigma_k^{-1}$ and $\|L_{22}\| - \sigma_{k+1}$ are thus cubic in σ_k^{-1} and σ_{k+1}, respectively. The derivation of the bounds illuminated the fact that one order of covergence is due to the rank-revealing pivoting in the first step; then, because of the pivoting in the first step, two more orders are achieved in the second step. In particular, the one order in the first step comes from the fact that we can bound the norms of R_{11}^{-1} and of R_{22} by σ_1^{-1} times a constant and σ_n times a constant, respectively, with the constants depending only on k and n. Our analysis assumes that $\Pi_1 = I$, that is, that pivoting is done only on the first step.

The algorithm can be continued beyond the first two steps, and we made some observations concerning the asymptotic convergence of individual elements. Assuming that Givens rotations are used to triangularize the matrix at each iteration, we saw that there were basically two effects that a Givens rotation can have on an off-diagonal element. When all of the off-diagonal elements are roughly the same size, then one type of effect dominates, and convergence is relatively slow. When, on the other hand, there are significant disparities in the sizes of off-diagonal elements, the other type of effect dominates, and the convergence is accelerated. This is the case, for example, when there are repeated singular values, providing more evidence as to how the pivoted QLP decomposition can converge so fast when singular values are close. We listed some of the common patterns convergence of off-diagonal elements.

We also showed that the asymptotic convergence of the diagonal element $r_{ss}^{(i)}$ is the square of the slowest asymptotic rate among all elements in row s and column s. We were then able to produce a bound on the asymptotic convergence rate of diagonal elements. Numerical examples illustrated the asymptotic convergence of diagonal and off-diagonal elements.

References

[1] T. F. Chan. Rank revealing QR factorizations. *Linear Algebra and Its Applications*, 88/89:67–82, 1987.

[2] S. Chandrasekaran and I.C.F. Ipsen. On rank-revealing factorisations. *SIAM Journal on Matrix Analysis and Applications*, 15(2):592–622, 1994.

[3] S. Chandrasekaran and I.C.F. Ipsen. Analysis of a QR algorithm for computing singular values. *SIAM Journal on Matrix Analysis and Applications*, 16(2):520–535, 1995.

[4] James W. Demmel. *Applied Numerical Linear Algebra.* SIAM, Philadelphia, 1997.

[5] L. V. Foster. Rank and null space calculations using matrix decomposition without column interchanges. *Linear Algebra and Its Applications*, 74:47–71, 1986.

[6] G. Golub and C. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, third edition, 1996.

[7] G. H. Golub. Numerical methods for solving least squares problems. *Numerische Mathematik*, 7:206–216, 1965.

[8] W. Kahan. Numerical linear algebra. *Canadian Math. Bulletin*, 9:757–801, 1966.

[9] R. Mathias and G. W. Stewart. A block QR algorithm and the singular value decomposition. *Linear Algebra and Its Applications*, 182:91–100, 1993.

[10] G. W. Stewart. *Introduction to Matrix Computations*. Academic Press, New York, 1973.

[11] G. W. Stewart. An updating algorithm for subspace tracking. *IEEE Transactions on Signal Processing*, 40:1535–1541, 1992.

[12] G. W. Stewart. Updating a rank-revealing ULV decomposition. *SIAM Journal on Matrix Analysis and Applications*, 14:494–499, 1993.

[13] G. W. Stewart. *Matrix Algorithms, Volume 1: Basic Decompositions*. SIAM, Philadelphia, 1998.

[14] G. W. Stewart. The QLP approximation to the singular value decomposition. *SIAM Journal on Scientific Computing*, 20(4):1336–1348, 1999.

[15] Lloyd N. Trefethen and David Bau, III. *Numerical Linear Algebra*. SIAM, Philadelphia, 1997.