

Chase and Trim: Earlier Deflation in Bidiagonal QR for the SVD

Eteri Byazrova^{a,*}, Yuri Morozov^a

^a*GraphMath, Coeur d'Alene, ID, USA*

Abstract

We introduce “Chase and Trim (CAT)” algorithm, an extension of the bidiagonal QR algorithm for computing singular values and singular vectors. CAT supplements conventional deflation with local diagonalization of selected interior 2×2 bidiagonal blocks. The resulting left and right rotations eliminate a superdiagonal coupling while introducing at most two cross-boundary fill-ins. When these fill-ins satisfy the local deflation tolerance, they are discarded, allowing the active block to split earlier. We derive algebraic conditions for rejecting infeasible candidates without constructing the rotations. Two inexpensive rejection tests screen candidates before invoking LAPACK DLASV2 for verification and application. We also show that the spectral norm of the perturbation discarded by an accepted CAT split does not exceed the tolerance permitted for conventional deflation. We compare standard bidiagonal QR with CAT using the same instrumented LAPACK DBDSQR implementation, with singular-vector accumulation enabled. Experiments cover 12 matrix families, 5 matrix orders from 64 to 1024, and 5 samples per family and order. Computational effort is measured by QR sweep positions and a weighted arithmetic-work model, including screening and singular-vector updates. CAT reduces weighted arithmetic work by 3.3% to 38.2% across the tested spectral families. Pooled reductions remained similar across the five tested matrix orders, ranging from 10.2% to 10.8%. The rejection screens eliminate approximately 99.95% of candidates before production verification, with no missed feasible opportunities observed. Singular-value accuracy, vector orthogonality, and reconstruction residuals remain comparable to conventional QR.

Keywords: bidiagonal QR, singular value decomposition, deflation, singular vectors, numerical linear algebra, Chase and Trim (CAT)

1. Introduction

The bidiagonal QR route for computing SVD of an $m \times n$ ($m \geq n$) matrix A proceeds as follows: (Watkins [1])

- (1) Reduce A to upper bidiagonal form by left and right Householder transformations. The resulting $m \times n$ bidiagonal matrix has zero rows below the leading $n \times n$ bidiagonal block; this square bidiagonal block B is retained for phase 2.
- (2) Choose a shift and apply a sweep of alternating right and left Givens rotations (LAPACK [4]) within the active block, chasing the temporary bulge from top to bottom and restoring bidiagonal form.
- (3) Scan the superdiagonal for conventional deflation opportunities between sweeps. Whenever a superdiagonal entry becomes sufficiently small, set it to zero and split the block at that position.
- (4) Repeat the process until every active block has size at most 2×2 and diagonalize directly

*Corresponding author

Email address: ebyazrova@aol.com (Eteri Byazrova)

If $m < n$, apply the algorithm to A^T instead, so the singular values remain unchanged, and the left and right singular vectors exchange roles.

Divide-and-conquer is an alternative method for computing the bidiagonal SVD. Bidiagonal QR nevertheless remains implemented in LAPACK DBDSQR, which supports computation and accumulation of the left and right singular vectors. (LAPACK [2]) CAT targets this QR method; comparison with divide-and-conquer is outside the present study.

We investigate if the same terminal 2×2 diagonalization can be applied at a non-negligible coupling. The transformation diagonalizes the 2×2 bidiagonal block (selected in Figure 1, left) and annihilates the selected superdiagonal entry (red). The transformation creates up to two new fill-in entries, shown in purple. If both new entries are below the accepted tolerance, they may be discarded and allow an early split.

We call this augmented bidiagonal QR strategy Chase and Trim (CAT).




Figure 1: Active CAT cut (left) and the cross-boundary fill-ins induced by the local two-sided rotations (right).

In between bulge chases, conventional deflations are applied as usual. After that, the remaining couplings are examined for active deflation feasibility. When a split is accepted, both rotations are applied and the child blocks continue recursively.

We compare the two algorithms for computing singular values and singular vectors.

- (1) Standard bidiagonal QR with only conventional deflation.
- (2) CAT in addition to (1).
 - Examines the rotation associated with the larger outer coupling.
 - Uses two inexpensive sufficient conditions to reject infeasible candidates.
 - Candidates surviving both screens are passed to LAPACK DLASV2 (LAPACK [6])

The resulting left and right rotations are used to verify both cross-boundary fill-ins. If both satisfy the tolerance, the rotations are applied and the block is split.

2. Active deflation by local SVD rotations

2.1. LBR product for active deflation

Let B be an $n \times n$ upper-bidiagonal matrix:

$$B = \begin{bmatrix} d_1 & e_1 & 0 & 0 & 0 & 0 \\ 0 & d_2 & e_2 & 0 & 0 & 0 \\ 0 & 0 & d_3 & e_3 & 0 & 0 \\ 0 & 0 & 0 & d_4 & e_4 & 0 \\ 0 & 0 & 0 & 0 & d_5 & e_5 \\ 0 & 0 & 0 & 0 & 0 & d_6 \end{bmatrix}$$

For the displayed example, $n = 6, k = 3$, the selected coupling is e_3 , between d_3 and d_4 .

Extract the selected 2×2 block:

$$K = \begin{bmatrix} d_3 & e_3 \\ 0 & d_4 \end{bmatrix}$$

Compute the local 2×2 SVD to obtain left and right rotations L_2 and R_2 . These rotations diagonalize K :

$$L_2 K R_2 = \begin{bmatrix} \sigma_3 & 0 \\ 0 & \sigma_4 \end{bmatrix}$$

The diagonal values may be signed, depending on the rotation convention. Embed L_2 and R_2 into $n \times n$ identity matrices L and R , acting only on the selected rows and columns. We now examine their effect on the full bidiagonal matrix, including neighboring couplings.

First, apply the right rotation R :

$$BR = \begin{bmatrix} d_1 & e_1 & 0 & 0 & 0 & 0 \\ 0 & d_2 & e_2 & 0 & 0 & 0 \\ 0 & 0 & d_3 & e_3 & 0 & 0 \\ 0 & 0 & 0 & d_4 & e_4 & 0 \\ 0 & 0 & 0 & 0 & d_5 & e_5 \\ 0 & 0 & 0 & 0 & 0 & d_6 \end{bmatrix} \times \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & c^R & -s^R & 0 & 0 \\ 0 & 0 & s^R & c^R & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} =$$

$$\begin{bmatrix} d_1 & e_1 & 0 & 0 & 0 & 0 \\ 0 & d_2 & c^R e_2 & -s^R e_2 & 0 & 0 \\ 0 & 0 & c^R d_3 + s^R e_3 & c^R e_3 - s^R d_3 & 0 & 0 \\ 0 & 0 & s^R d_4 & c^R d_4 & e_4 & 0 \\ 0 & 0 & 0 & 0 & d_5 & e_5 \\ 0 & 0 & 0 & 0 & 0 & d_6 \end{bmatrix}$$

Next, apply the left rotation L :

$$LBR = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & c^L & s^L & 0 & 0 \\ 0 & 0 & -s^L & c^L & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} d_1 & e_1 & 0 & 0 & 0 & 0 \\ 0 & d_2 & c^R e_2 & -s^R e_2 & 0 & 0 \\ 0 & 0 & c^R d_3 + s^R e_3 & c^R e_3 - s^R d_3 & 0 & 0 \\ 0 & 0 & s^R d_4 & c^R d_4 & e_4 & 0 \\ 0 & 0 & 0 & 0 & d_5 & e_5 \\ 0 & 0 & 0 & 0 & 0 & d_6 \end{bmatrix} =$$

$$\begin{bmatrix} d_1 & e_1 & 0 & 0 & 0 & 0 \\ 0 & d_2 & c^R e_2 & -s^R e_2 & 0 & 0 \\ 0 & 0 & c^L (c^R d_3 + s^R e_3) + s^L s^R d_4 & c^L (c^R e_3 - s^R d_3) + s^L c^R d_4 & s^L e_4 & 0 \\ 0 & 0 & -s^L (c^R d_3 + s^R e_3) + c^L s^R d_4 & -s^L (c^R e_3 - s^R d_3) + c^L c^R d_4 & c^L e_4 & 0 \\ 0 & 0 & 0 & 0 & d_5 & e_5 \\ 0 & 0 & 0 & 0 & 0 & d_6 \end{bmatrix}$$

By construction, the local SVD rotations annihilate both off-diagonal entries of the selected block: $-s^L (c^R d_3 + s^R e_3) + c^L s^R d_4 = 0$ $c^L (c^R e_3 - s^R d_3) + s^L c^R d_4 = 0$.

$$LBR = \begin{bmatrix} d_1 & e_1 & 0 & 0 & 0 & 0 \\ 0 & d_2 & c^R e_2 & -s^R e_2 & 0 & 0 \\ 0 & 0 & \sigma_3 & 0 & s^L e_4 & 0 \\ 0 & 0 & 0 & \sigma_4 & c^L e_4 & 0 \\ 0 & 0 & 0 & 0 & d_5 & e_5 \\ 0 & 0 & 0 & 0 & 0 & d_6 \end{bmatrix}$$

The selected coupling is eliminated, but two cross-boundary fill-ins remain:

- left fill $= -s^R e_2$

- right fill = $s^L e_4$.

The same fill-ins result whether we compute $L(B R)$ or $(L B)R$.

Early deflation is feasible when both fill-in magnitudes satisfy the local DBDSQR deflation tolerance τ (LAPACK [2]) :

- $|-s^R e_2| \leq \tau$
- $|s^L e_4| \leq \tau$.

2.2. From fill-in to CAT inequalities

Consider the same upper-bidiagonal matrix:

$$B = \begin{bmatrix} d_1 & e_1 & 0 & 0 & 0 & 0 \\ 0 & d_2 & e_2 & 0 & 0 & 0 \\ 0 & 0 & d_3 & e_3 & 0 & 0 \\ 0 & 0 & 0 & d_4 & e_4 & 0 \\ 0 & 0 & 0 & 0 & d_5 & e_5 \\ 0 & 0 & 0 & 0 & 0 & d_6 \end{bmatrix}$$

Suppose we seek to rotate and eliminate the coupling e_3 , between d_3 and d_4 . Denote

- $L = |e_2|$
- $R = |e_4|$

(L and R denote neighboring coupling magnitudes, not the rotation matrices of the preceding section)

- τ = directional deflation tolerance for the selected index (LAPACK [2]).

Define μ_k as the recurrence value for deflation at e_k . If the chase proceeds left-to-right:

- $\mu_1 = |d_1|$
- $\mu_{k+1} = \frac{|d_{k+1}|\mu_k}{(\mu_k + |e_k|)}$
- (deflation threshold at e_k) or $\tau_k = \text{tol } \mu_k$

If the chase proceeds right-to-left inside a block with h columns:

- $\mu_h = |d_h|$
- $\mu_k = \frac{|d_k|\mu_{k+1}}{(\mu_{k+1} + |e_k|)}$
- (deflation threshold at e_k) or $\tau_k = \text{tol } \mu_{k+1}$

(note that $\mu_k \geq 0$ for all k and, for the cut at e_k , we denote τ_k by τ)

The two rotations create the fill-ins identified earlier:

- the right rotation creates a left fill-in of magnitude $L|s^R|$
- the left rotation creates a right fill-in of magnitude $R|s^L|$

For active deflation, both magnitudes must be no larger than τ .

Since $L(BR) = (LB)R$, the left and right rotations may be constructed in either order.

Suppose we start with the right rotation, which is

- constructed from diagonalization of $K^T K$
- allows deflation only if $L|s^R| \leq \tau$

For the selected e, define quantities of local

$$K^T K = \begin{bmatrix} d_3^2 & d_3 e_3 \\ d_3 e_3 & e_3^2 + d_4^2 \end{bmatrix}$$

- $A^R = d_3^2$
- $C^R = e_3^2 + d_4^2$
- $q^R = \text{abs}(d_3 e_3)$

(for a nontrivial rotation considered below, $q^R > 0$)

- $\Delta^R = \text{abs}(A^R - C^R)$

(if $\Delta^R = 0$, $\text{abs}(\theta^R) = \frac{\pi}{4}$ and feasibility is decided directly from $\frac{L}{\sqrt{2}} \leq \tau$)

Let θ^R be the angle of the smaller right rotation, so

- $s^R = \sin \theta^R$
- $c^R = \cos \theta^R$

the rotation angle is computed as

$$\text{abs}(\tan(2\theta^R)) = \frac{2q^R}{\Delta^R}$$

From

$$\tan(2\theta^R) = \frac{2 \tan \theta^R}{1 - (\tan \theta^R)^2}$$

we obtain

$$\begin{aligned} \frac{\Delta^R}{2q^R} &= \frac{1 - (\tan \theta^R)^2}{2 \text{abs}(\tan \theta^R)} \\ &= \frac{1}{2 \text{abs}(\tan \theta^R)} - \frac{\text{abs}(\tan \theta^R)}{2} \end{aligned}$$

For the shorter rotation, $\text{abs}(\theta^R) \leq \frac{\pi}{4}$ Therefore $\text{abs}(\tan \theta^R)$ increases with $\text{abs}(\theta^R)$, while the Jacobi ratio above decreases.

For feasibility testing only, define

- $s = \text{abs}(s^R)$

Since s also increases with $\text{abs}(\theta^R)$, the Jacobi ratio decreases as s increases.

Work backward from the limit of feasible rotation assuming $L > \tau > 0$; otherwise this fill-in is automatically acceptable when $L \leq \tau$

$$s = \frac{\tau}{L}$$

Since $s = \text{abs}(\sin \theta^R)$,

$$\text{abs}(\tan \theta^R) = \frac{s}{\sqrt{1-s^2}},$$

and therefore

$$\frac{\Delta^R}{2q^R} = \frac{1-2s^2}{(2s\sqrt{1-s^2})}$$

Substituting $s = \frac{\tau}{L}$ into the above equation gives

$$\Delta^R \tau \sqrt{L^2 - \tau^2} = q^R (L^2 - 2\tau^2).$$

Because the Jacobi ratio decreases as s increases, the right rotation is feasible when

$$\Delta^R \tau \sqrt{L^2 - \tau^2} \geq q^R (L^2 - 2\tau^2)$$

2.3. Rejection conditions

The left fill is too large for deflation when

$$\Delta^R \tau \sqrt{L^2 - \tau^2} < q^R (L^2 - 2\tau^2)$$

From this exact rejection condition, we derive two early conditions, both for $L \geq 2\tau$

$$(1) \quad \sqrt{L^2 - \tau^2} < L$$

$$\text{If } L \geq 2\tau, L^2 - 2\tau^2 \geq L(L - \tau).$$

To obtain a simpler rejection condition, require an upper bound on the left side of

$$\Delta^R \tau \sqrt{L^2 - \tau^2} < q^R (L^2 - 2\tau^2)$$

to be smaller than a lower bound on its right side:

$$\Delta^R \tau L < q^R L(L - \tau)$$

and, since $L > 0$,

$$\Delta^R \tau < q^R (L - \tau).$$

(2) A simpler rejection condition follows from the first. If $L \geq 2\tau$, then $L - \tau \geq \tau$. Therefore, if $\Delta^R \leq q^R$,

$$\Delta^R \tau \leq q^R \tau \leq q^R (L - \tau).$$

Thus we can also reject when $L \geq 2\tau$ and $\Delta^R \leq q^R$.

The simpler condition is derived from the stronger one. Although derived second, it is tested first because it requires only one addition. The stronger condition requires a subtraction and two multiplications. It is evaluated only when $\Delta^R > q^R$.

The stronger condition

$$\Delta^R \tau < q^R (L - \tau)$$

needs no separate $L \geq 2\tau$ comparison. It is evaluated only when $\Delta^R > q^R$, so

$$q^R \tau \leq \Delta^R \tau < q^R (L - \tau).$$

Dividing by q^R gives $\tau < L - \tau$, hence $L > 2\tau$

($q^R > 0$ was established previously)

2.4. Evaluation order

Candidate cuts are examined from the destination of the QR chase backward. This prioritizes the region where deflations are more likely to occur. The first feasible cut is accepted. DBDSQR then resumes block processing, checking for further deflations before performing another QR sweep.

For each candidate cut, Rejection CAT proceeds as follows:

- Evaluate the rotation associated with the larger outer coupling.
- Apply the inexpensive rejection condition first.
- Apply the stronger rejection condition only if the first test fails to reject.
- Pass candidates surviving both tests directly to LAPACK DLASV2.

The two rejection conditions, in order, are:

- (1) Reject if $L \geq 2\tau$ and $\Delta^R \leq q^R$.
- (2) Reject if $\Delta^R \tau < q^R(L - \tau)$.
- (3) For surviving candidates, compute both rotations using LAPACK DLASV2 (LAPACK [6]).
- (4) Verify the two cross-boundary fill-ins against τ . Reject the candidate immediately if either fill-in exceeds τ .
- (5) If both fill-ins satisfy τ , apply both rotations, update the singular vectors, and split the active block.

Numerical errors in screening may affect efficiency, but not split validity, which is checked by production verification.

2.5. CAT algorithm: pseudocode

Input: active upper-bidiagonal block $B[\text{l..r}]$, tolerance τ_k at each cut k , and QR chase direction idir. After conventional deflation, examine the surviving couplings in the direction opposite to the QR chase:

- If idir = 1 (left-to-right chase), scan $k = r - 1, \dots, 1$.
- If idir = 2 (right-to-left chase), scan $k = 1, \dots, r - 1$.

At each candidate cut, $b = e_k$. Apply the following tests in scan order.

1. Obtain the local quantities

$$\begin{aligned}
 d_1 &= B[k, k] \\
 d_2 &= B[k + 1, k + 1] \\
 b &= B[k, k + 1] \\
 L &= |e_{k-1}| \text{ if } k > l, \text{ otherwise } 0 \\
 R &= |e_{k+1}| \text{ if } k + 1 < r, \text{ otherwise } 0 \\
 \tau &= \tau_k, \quad M = \max(L, R)
 \end{aligned}$$

2. If $M \leq \tau$, go to step 6.
3. Choose the less permissive rotation and construct its Gram matrix quantities.
If $L \geq R$, use $K^T K$:

$$\begin{aligned}
 A &= d_1^2 \\
 C &= b^2 + d_2^2 \\
 q &= d_1 b
 \end{aligned}$$

Otherwise, use KK^T :

$$A = d_1^2 + b^2$$

$$C = d_2^2$$

$$q = bd_2$$

$$\Delta = A - C$$

4. First rejection screen:

If $|\Delta| \leq |q|$:

If $M \geq 2\tau$, reject this cut and continue to the next k .

Otherwise, proceed directly to step 6.

5. Second rejection screen:

If $|\Delta| > |q|$, compute:

$$M - \tau$$

$$|\Delta|\tau$$

$$|q|(M - \tau)$$

If $|\Delta|\tau < |q|(M - \tau)$, reject this cut and continue to the next k .

6. Compute the local 2×2 SVD using LAPACK DLASV2 (LAPACK [6]).

Obtain right rotation (c^R, s^R) , left rotation (c^L, s^L) , and signed diagonal values.

Use the same ordering convention as the production algorithm.

7. Verify the cross-boundary fills.

If $k > l$ and $|s^R e_{k-1}| > \tau$, reject this cut and continue.

If $k + 1 < r$ and $|s^L e_{k+1}| > \tau$, reject this cut and continue.

8. If the cut is accepted

Scale existing outer couplings by c^R and c^L where present

Replace the two local diagonal entries with the returned signed values

Set $e_k = 0$

Accumulate the right rotation in V^T and the left rotation in U

Return to bidiagonal QR to process the resulting blocks

for conventional deflation

if no conventional deflation occurs, active deflation

if no active deflation occurs, new bulge chase

9. If every candidate is rejected, perform the ordinary bidiagonal QR sweep.

2.6. Why do the two rejection screens miss so little?

The first two tests reject a candidate before the exact fill condition is evaluated.

In Figure 2, the green curve shows the exact rejection boundary, while the blue curve shows the inexpensive rejection boundary from the second screen.




Figure 2: The exact rejection boundary (green) and the inexpensive sufficient rejection boundary (blue), plotted as $r = \Delta/q$ against $x = M/\tau$.

For the less permissive rotation, define

- $M = \max(L, R)$
- $r = \frac{\Delta}{q}$
- $x = \frac{M}{\tau}$

Recall that $M \geq 2\tau$, so $x \geq 2$ is part of our rejection setting.

Substituting r and $M = \tau x$ into the exact rejection condition

$$\Delta\tau \sqrt{M^2 - \tau^2} < q(M^2 - 2\tau^2)$$

results in $r \sqrt{x^2 - 1} < x^2 - 2$, or, equivalently,

$$r < \frac{x^2 - 2}{\sqrt{x^2 - 1}}$$

Substituting r and $M = \tau x$ into the second rejection screen

$$\Delta\tau < q(M - \tau)$$

results in $r < x - 1$.

Thus, for $x \geq 2$, the blue boundary lies below the green boundary. The region between them shows the small gap that corresponds to possible candidates rejected by the exact test but not by the inexpensive screen.

To examine the gap for large x , rewrite the exact rejection boundary as

$$\frac{x^2 - 2}{\sqrt{x^2 - 1}} = x \left(1 - \frac{2}{x^2}\right) \left(1 - \frac{1}{x^2}\right)^{-1/2}$$

Define

$$u = \frac{1}{x^2}, \quad v = 1 - u$$

$$f(u) = \left(1 - \frac{1}{x^2}\right)^{-1/2} = (1 - u)^{-0.5} = v^{-0.5}$$

Since $\frac{dv}{du} = -1$

$$f'(u) = \left(\frac{1}{2}\right) v^{-1.5}$$

$$f''(u) = \left(\frac{3}{4}\right) v^{-2.5}$$

At $u = 0, v = 1$, so

$$f(0) = 1, f'(0) = \frac{1}{2}, f''(0) = \frac{3}{4}$$

The Taylor expansion around $u = 0$ is

$$\begin{aligned} f(u) &= f(0) + f'(0)u + \frac{f''(0)u^2}{2!} + \dots \\ &= 1 + \frac{u}{2} + \frac{3u^2}{8} + \frac{5u^3}{16} + \dots \end{aligned}$$

Multiplying by $(1 - 2u)$, substituting $u = \frac{1}{x^2}$, and multiplying by x gives

$$\frac{x^2 - 2}{\sqrt{x^2 - 1}} = x - \frac{3}{2x} - \frac{5}{8x^3} - \dots$$

Subtracting the inexpensive rejection boundary $x - 1$ gives

$$G(x) = 1 - \frac{3}{2x} - \frac{5}{8x^3} - \dots$$

The Taylor series converges here because $u = \frac{1}{x^2} < 1$. Thus, the absolute gap approaches 1 from below, while the gap relative to x approaches 0. Its leading relative term is $\frac{1}{x} = \frac{\tau}{M}$.

3. Singular-vector accumulation

Accumulation of the Householder reflectors from phase 1 is excluded from this comparison.

The phase-2 singular-vector matrices U and V^T are initialized to identity $\mathcal{I}_n$. During phase 2 one left and one right Givens rotation is applied by:

- Each QR chase position
- Terminal 2×2 diagonalization
- An accepted CAT split

Conventional deflation requires no additional singular-vector rotation.

For an active bidiagonal block of order h , the affected panels are:

- $n \times h$ columns of U
- $h \times n$ rows of V^T




Figure 3: Singular-vector accumulation: two columns of U (top) and two rows of V^T (bottom) updated by a rotation. Blue dots identify the updated vectors.

Example: $n = 9, h = 5$, active indices 3...7 (Figure 3). Filled dots identify the active panel; blue marks the two vectors updated by one rotation.

Each rotation updates two vectors of length n . For each entry pair, four multiplications and two additions or subtractions are required.

- One singular-vector rotation: $6n$
- One paired left and right update: $2(6n) = 12n$
- One QR sweep on an active block of order h : $12n(h-1)$
- One accepted CAT split: $12n$
- One terminal 2×2 diagonalization: $12n$

These updates are performed by the modified LAPACK DBDSQR production function.

4. Spectral-norm bound for accepted splits

The same LAPACK DLASV2 rotations diagonalize a terminal 2×2 block or an interior one. An interior rotation differs only in the two cross-boundary fill-ins that CAT discards.

Let $B' = LBR$ be the matrix after the rotations and $\tilde{B}$ the matrix after trimming. Then $B' = \tilde{B} + E$. The discarded perturbation in the four affected coordinates is:

$$E = \begin{bmatrix} 0 & 0 & -s^R e_2 & 0 \\ 0 & 0 & 0 & s^L e_4 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$E^T E = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -s^R e_2 & 0 & 0 & 0 \\ 0 & s^L e_4 & 0 & 0 \end{bmatrix} \times \begin{bmatrix} 0 & 0 & -s^R e_2 & 0 \\ 0 & 0 & 0 & s^L e_4 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & (s^R e_2)^2 & 0 \\ 0 & 0 & 0 & (s^L e_4)^2 \end{bmatrix}$$

The two fill-ins occupy different rows and columns, so $E^T E$ is diagonal.

$$\lambda_{\max}(E^T E) = \max\left((s^R e_2)^2, (s^L e_4)^2\right).$$

$$\|E\|_2 = \sqrt{\lambda_{\max}(E^T E)} = \max(|s^R e_2|, |s^L e_4|).$$

CAT accepts only if $|s^R e_2| \leq \tau$ and $|s^L e_4| \leq \tau$. Therefore $\|E\|_2 \leq \tau$.

Conventional deflation discards only e_3 :

$$E_p = \begin{bmatrix} 0 & e_3 \\ 0 & 0 \end{bmatrix}$$

$$E_p^T E_p = \begin{bmatrix} 0 & 0 \\ 0 & e_3^2 \end{bmatrix}$$

$$\|E_p\|_2 = \sqrt{\lambda_{\max}(E_p^T E_p)} = |e_3| \leq \tau.$$

Thus conventional and active splits enforce the same spectral-norm bound on each discarded perturbation.

5. Numerical experiments and accounting

5.1. Computational methods

Both methods begin with Householder bidiagonalization of the original rectangular matrix A . The identical upper-bidiagonal matrix B is supplied to both phase-2 algorithms.

Both use the same instrumented production LAPACK DBDSQR implementation for:

- Shift selection and bidiagonal QR bulge chasing
- Conventional deflation
- Terminal 2×2 diagonalization using LAPACK DLASV2
- Singular-vector accumulation

CAT adds only:

- Screening for additional deflation opportunities between QR sweeps
- Applying the same LAPACK DLASV2 local diagonalization to selected nonterminal 2×2 blocks
- Verifying and discarding negligible fill-ins, updating singular vectors, and splitting the active block

Computational effort is evaluated for phase 2 only. The common Householder bidiagonalization is excluded from the comparison.

The weighted-work model uses structure-aware arithmetic counts for standard LAPACK bidiagonal QR iteration. The local 2×2 diagonalization and singular-vector rotation costs are modeled from LAPACK DLASV2 and BLAS DROT. Shift computation is charged using the branch-dependent work recorded by the instrumented DBDSQR routine.

We report QR sweep positions and weighted arithmetic work, not execution time.

5.2. Matrix families used for testing

The matrix families are listed in Table 1.

Table 1: Matrix families used for testing. Matrix dimensions: $2n \times n$.

Matrix family	Construction for the experiment	Adjacent σ relation, if applicable	Related generator, if applicable
Power law, $\alpha = 0.5$ Specified singular values with power-law decay	Set $\sigma_i = i^{-0.5}$, $i = 1, \dots, n$. Construct a $2n \times n$ diagonal matrix Σ containing the specified singular values. Construct U and V as products of $4 \max(m, n)$ plane rotations, with angles in $[0.2, 1.35]$ radians. Construct $A = U\Sigma V^T$.	$\sigma_{k+1} = \sigma_k \left(\frac{k}{k+1}\right)^{0.5}$	cf. LAPACK [8]
$\alpha = 1.0$	Set $\sigma_i = i^{-1.0}$, $i = 1, \dots, n$. Otherwise as above.	$\sigma_{k+1} = \sigma_k \left(\frac{k}{k+1}\right)^{1.0}$	cf. LAPACK [8]
$\alpha = 2.0$	Set $\sigma_i = i^{-2.0}$, $i = 1, \dots, n$. Otherwise as above.	$\sigma_{k+1} = \sigma_k \left(\frac{k}{k+1}\right)^{2.0}$	cf. LAPACK [8]
$\alpha = 3.0$	Set $\sigma_i = i^{-3.0}$, $i = 1, \dots, n$. Otherwise as above.	$\sigma_{k+1} = \sigma_k \left(\frac{k}{k+1}\right)^{3.0}$	cf. LAPACK [8]
IID noise Independent uniformly distributed entries	Sample each entry independently from $[-1.0, 1.0]$.		
Uniform $[1.0, 10.0]$ Uniformly sampled singular values	Sample singular values from $[1.0, 10.0]$. Construct a $2n \times n$ diagonal matrix Σ containing the specified singular values. Construct U and V as products of $4 \max(m, n)$ plane rotations, with angles in $[0.2, 1.35]$ radians. Construct $A = U\Sigma V^T$.		cf. LAPACK [8]
$[1.0, 100.0]$	Sample singular values from $[1.0, 100.0]$. Otherwise as above.		cf. LAPACK [8]
$[1.0, 1000.0]$	Sample singular values from $[1.0, 1000.0]$. Otherwise as above.		cf. LAPACK [8]
$[1.0, 10000.0]$	Sample singular values from $[1.0, 10000.0]$. Otherwise as above.		cf. LAPACK [8]
Linear Linearly spaced singular values	Set $\sigma_i = 0.001 + (i-1)h$, $i = 1, \dots, n$. $h = (1000.0 - 0.001)/(n-1)$. Construct a $2n \times n$ diagonal matrix Σ containing the specified singular values. Construct U and V as products of $4 \max(m, n)$ plane rotations, with angles in $[0.2, 1.35]$ radians. Construct $A = U\Sigma V^T$.	$\sigma_{k+1} = \sigma_k + h$	cf. LAPACK [8]
Geometric Geometrically spaced singular values	Set $\sigma_i = 0.001 \times r^{i-1}$, $i = 1, \dots, n$. $r = (10^6)^{1/(n-1)}$. Construct a $2n \times n$ diagonal matrix Σ containing the specified singular values. Construct U and V as products of $4 \max(m, n)$ plane rotations, with angles in $[0.2, 1.35]$ radians. Construct $A = U\Sigma V^T$.	$\sigma_{k+1} = r\sigma_k$	cf. LAPACK [8]
Tightly clustered $[1.0, 1.001]$ Tightly clustered	Sample singular values from $[1.0, 1.001]$. Construct a $2n \times n$ diagonal matrix Σ containing the specified singular values. Construct U and V as products of $4 \max(m, n)$ plane rotations, with angles in $[0.2, 1.35]$ radians. Construct $A = U\Sigma V^T$.		cf. LAPACK [8]

The spectral distributions used in these experiments are similar to those employed in the LAPACK solver testing infrastructure (Demmel et al. [9]).

For each matrix family, 5 deterministic matrices are tested at $n = 64, 128, 256, 512, 1024$. The random seeds are 1, 2, 3, 4, 5.

5.3. Reported quantities

Tables 2–5 define the reported quantities: computational effort, numerical accuracy, rejection screens, and production verification.

Table 2: Computational effort.

Quantity	Formula	Reported
Sweep positions	$\sum_i (h - 1)$, over all active blocks and QR sweeps	Average across samples
Weighted arithmetic work W		Average across samples
CAT test work		Average across samples
Accepted CAT work		Average across samples

Conventional deflations are ordinary bidiagonal QR splits. Active CAT deflations are the splits accepted by the production library. Terminal 2×2 diagonalizations are excluded from both counts.

$$\text{Active CAT fraction} = \frac{\text{active CAT}}{\text{active CAT} + \text{conventional}}$$

For a parent block of size m and children of sizes n_l and n_r ,

$$\text{smaller-child ratio} = \frac{\min(n_l, n_r)}{m}$$

The ratio lies between $\frac{1}{m}$ and $\frac{1}{2}$

We report the mean smaller-child ratio separately for

- active CAT splits
- conventional splits in CAT
- conventional splits in standard QR.

Counts are summed across samples; mean balances are pooled across individual splits. The counts do not generally sum to $5(n - 1)$, because terminal 2×2 blocks are excluded.

Table 3: Numerical accuracy and singular-vector quality.

Quantity	Formula	Reported
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	Maximum across all matrix sizes and samples
Max scaled σ error (floor = 1)	$\max_i \left(\frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)} \right)$	Maximum across all matrix sizes and samples
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \left(\frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* } \right)$	Maximum across all matrix sizes and samples
Left singular-vector orthogonality	$\ U^T U - I\ _F$	Maximum across all matrix sizes and samples
Right singular-vector orthogonality	$\ V^T V - I\ _F$	Maximum across all matrix sizes and samples
Relative SVD residual	$\ B - U \Sigma V^T\ _F / \ B\ _F$	Maximum across all matrix sizes and samples

The reference singular values σ_i^* are computed by LAPACK DGESVD from the original rectangular matrix A . The orthogonality and reconstruction-residual measures follow the standard forms used in LAPACK eigensolver testing (Demmel et al. [9]), adapted here to the SVD.

Table 4: CAT rejection screens.

Quantity	Reported
Total CAT candidate inspections	Sum across samples
First-test rejections	Sum across samples
Second-test rejections	Sum across samples

Table 5: Production verification.

Quantity	Reported
Missed CAT opportunities	Sum across samples
Sent to production	Sum across samples
Rejected by production	Sum across samples

5.4. Weighted-work accounting

Table 6 gives the relative operation weights, and Table 7 the arithmetic model for each operation.

Table 6: Relative weights of arithmetic operations.

Operation	Weight
Floating-point addition	1
Floating-point subtraction	1
Floating-point multiplication	1
Floating-point division	4
Square root	4
hypot(x, y)	8

Table 7: Weighted-work arithmetic models.

Quantity		Operation model	Weighted work	Source
Singular-vector rotation	rota-	One rotation applied to U or V	$6n$	LAPACK [7]
Conventional convergence scan	conver-	Initial test + up to $(h - 1)$ coupling tests (stops at first deflation)	$1 + 7(h - 1)$ (full scan only)	LAPACK [2]
Shift computation		Shift selection + optional DLAS2 computation + optional negligibility test (cost recorded by production)	7–53 (branch-dependent)	LAPACK [2], [3]
Zero-shift QR sweep, bidiagonal updates	bidiagonal	Two Givens + simplified updates per chase position + endpoint correction + shift selection & computation	$36(h - 1) + 2 + \text{shift work}$	LAPACK [2], [4]
Shifted bidiagonal QR sweep, bidiagonal updates		Two Givens + full updates per chase position + net boundary correction $(7 - 2)$ + shift selection & computation	$48(h - 1) + 5 + \text{shift work}$	LAPACK [2], [4]
Singular-vector updates during QR sweep		One U rotation + one V rotation per chase position	$2(6)n(h - 1)$	LAPACK [5]
CAT rejection, test 1		Three squares + one diagonal sum + one off-diagonal product + one diagonal difference + tolerance doubling	$3 + 1 + 1 + 1 + 1 = 7$	This work
CAT rejection, test 2		Three squares + one diagonal sum + one off-diagonal product + one diagonal difference + margin subtraction + two products	$3 + 1 + 1 + 1 + 1 + 2 = 9$	This work

Quantity	Operation model	Weighted work	Source
Production local 2×2 SVD	DLASV2 computes candidate rotations (no singular-vector updates)	76	LAPACK [6]
Terminal 2×2 diagonalization	DLASV2 computation + rotations applied to U and V	$76 + 2(6)n$	LAPACK [6], [7]
Accepted CAT, end-point	DLASV2 computation + one outer-fill verification + one surviving-coupling update + rotations applied to U and V	$76 + 1 + 1 + 2(6)n$	LAPACK [6], [7]
Accepted CAT, interior	DLASV2 computation + two outer-fill verifications + two surviving-coupling updates + rotations applied to U and V	$76 + 2 + 2 + 2(6)n$	LAPACK [6], [7]

Here n is the full bidiagonal order and h is the current active-block size. The shift cost is supplied by the instrumented DBDSQR algorithm for each QR sweep.

CAT work is charged as follows:

- Rejection-screen tests
- Local DLASV2 computation
- Fill-in verification
- Surviving-coupling updates
- Singular-vector rotations

Comparisons, assignments, sign changes, and diagnostic-only computations are excluded.

The cost of bidiagonalization is also excluded.

6. Numerical results

The full family-by-family sweep-position, weighted-work, deflation, screening, and numerical-accuracy tables are preserved in Appendix A. Table 8 gives the weighted-work reduction for all five tested matrix orders and pooled reductions over orders and families.

Table 8: Percentage reduction in weighted arithmetic work, CAT compared with standard QR, including screening, accepted-CAT rotations, and singular-vector updates. Pooled percentages are ratios of summed work reductions to summed baseline work.

Family	$n = 64$	$n = 128$	$n = 256$	$n = 512$	$n = 1024$	Pooled over n
Power-law $\alpha = 0.5$	7.51%	9.30%	7.13%	6.91%	7.06%	7.05%
Power-law $\alpha = 1.0$	9.37%	10.59%	10.44%	10.22%	9.89%	9.93%
Power-law $\alpha = 2.0$	13.37%	15.77%	15.56%	14.65%	15.17%	15.12%
Power-law $\alpha = 3.0$	23.20%	22.47%	23.07%	23.80%	22.53%	22.68%
IID rectangular noise	4.28%	2.81%	5.03%	4.61%	3.14%	3.33%
Uniform [1, 10]	4.22%	4.85%	4.74%	5.84%	6.48%	6.38%
Uniform [1, 100]	4.30%	3.94%	5.14%	4.76%	5.88%	5.74%
Uniform [1, 1000]	4.42%	4.72%	4.47%	5.56%	5.60%	5.58%
Uniform [1, 10000]	3.78%	3.88%	3.60%	4.83%	5.56%	5.45%
Linearly spaced [0.001, 1000]	3.43%	4.04%	5.08%	5.75%	6.01%	5.97%
Geometric [0.001, 1000]	43.93%	43.07%	42.44%	40.38%	37.84%	38.21%
Tightly clustered [1, 1.001]	4.09%	7.67%	5.61%	5.15%	5.66%	5.61%
Pooled over families	10.25%	10.82%	10.83%	10.82%	10.66%	10.68%

7. Spectral separation and CAT deflation

The observed CAT benefit varies substantially across spectral families and is most pronounced in the geometric group.

We examine how singular-value separation enters the local deflation condition. We consider $|e_2| \geq |e_4|$, the commonly observed ordering, where the right rotation then gives the more restrictive fill-in condition.

At the selected coupling, the local bidiagonal block is K

$$= \begin{bmatrix} d_3 & e_3 \\ 0 & d_4 \end{bmatrix}$$

The right rotation diagonalizes local $K^T K$

$$= \begin{bmatrix} d_3^2 & d_3 e_3 \\ d_3 e_3 & e_3^2 + d_4^2 \end{bmatrix}$$

The left rotation diagonalizes local $K K^T$

$$= \begin{bmatrix} d_3^2 + e_3^2 & e_3 d_4 \\ e_3 d_4 & d_4^2 \end{bmatrix}$$

The corresponding diagonal differences and off-diagonal magnitudes are:

$$\Delta^R = |d_3^2 - d_4^2 - e_3^2|$$

$$q^R = |d_3 e_3|$$

$$\Delta^L = |d_3^2 + e_3^2 - d_4^2|$$

$$q^L = |e_3 d_4|$$

Recall the exact right-rotation feasibility condition:

$$\Delta^R \tau \sqrt{L^2 - \tau^2} \geq q^R (L^2 - 2\tau^2)$$

Substituting Δ^R, q^R , and $L = |e_2|$ gives:

$$|d_3^2 - d_4^2 - e_3^2| \tau \sqrt{e_2^2 - \tau^2} \geq |d_3 e_3| (e_2^2 - 2\tau^2).$$

For $|e_2| \gg \tau$, this condition is approximately:

$$\frac{|d_3^2 - d_4^2 - e_3^2|}{|d_3 e_3|} \gtrsim \frac{|e_2|}{(\tau)}$$

or

$$\tau |d_3^2 - d_4^2 - e_3^2| \geq |d_3 e_3 e_2|$$

Entries near convergence:

- e_3 gets smaller and e_3^2 becomes negligible
- diagonal magnitudes approach singular values
- diagonal entries are arranged approximately in spectral order

Let $M = \max(\sigma_3, \sigma_4)$ and $m = \min(\sigma_3, \sigma_4)$. The corresponding approximate condition for the right rotation becomes

$$|e_2 e_3| \leq \tau \frac{M^2 - m^2}{M}$$

The left rotation must also satisfy its similar condition.

7.1. Matrices with linear σ spacing

For consecutive values, $m = M - h$, where $h > 0$.

$$M^2 - m^2 = 2Mh - h^2.$$

Therefore,

$$\begin{aligned} |e_2 e_3| &\leq \tau \frac{2Mh - h^2}{M} \\ &= \tau \left(2h - \frac{h^2}{M} \right). \end{aligned}$$

The linear factor $2h - \frac{h^2}{M}$ is approximately constant when $M \gg h$.

7.2. Matrices with geometric σ spacing

Let $\rho = \frac{1}{r}$, where $r > 1$ is the ratio used in the tables. For consecutive values, $m = \rho M$, where $0 < \rho < 1$.

$$M^2 - m^2 = (1 - \rho^2) M^2.$$

Therefore,

$$\begin{aligned} |e_2 e_3| &\leq \tau \frac{(1 - \rho^2) M^2}{M} \\ &= \tau (1 - \rho^2) M. \end{aligned}$$

The geometric factor $(1 - \rho^2) M$ scales with the larger singular value and may facilitate interior deflations, consistent with the experiments.

The experiment also shows that geometric spectra produce more interior splits under both conventional QR deflation and active CAT deflation.

8. Conclusions

8.1. Algorithm

The Chase and Trim (CAT) method provides additional deflation opportunities for SVD, augmenting conventional bidiagonal QR with local 2×2 SVD diagonalizations. Conventional deflation opportunities are evaluated and applied first. The remaining superdiagonal entries are then examined for active deflation. When a feasible candidate is found, the selected coupling is eliminated, the local diagonal entries and singular vectors are updated, and the active block is split into two smaller blocks.

Deflation opportunities are favored by small superdiagonal entries and large differences between neighboring diagonal entries. These conditions become more common as QR iteration approaches convergence.

We show that the spectral norm of the perturbation introduced by active deflation does not exceed that permitted by conventional passive deflation.

We present efficient rejection tests that eliminate unsuitable candidates at low cost. Candidates that pass these tests are submitted to the production function, which makes the final feasibility decision and, if successful, applies the transformation to all three matrices without explicitly forming $K^T K$.

8.2. Reduction in computational work

After inclusion of candidate testing and accepted rotations, CAT reduces weighted arithmetic work, ranging from 3.3% for IID rectangular noise to 38.2% for Geometric singular values $[0.001, 1000.0]$, pooled across samples and matrix sizes.

8.3. Dependence on spectral structure

The effectiveness of CAT depends strongly on the singular-value distribution. The largest work reductions occur for Geometric singular values $[0.001, 1000.0]$, and the smallest for IID rectangular noise. This difference may be explained in part by split balance. The mean smaller-child-to-parent ratio for active CAT splits, pooled over individual accepted splits across samples and matrix sizes, was 0.013 for the latter family, compared with 0.146 for the former.

8.4. Dependence on matrix size

Pooled across the twelve matrix families, the relative reduction in weighted arithmetic work ranges from 10.2% to 10.8% over the five tested matrix orders. The variation between matrix families is substantially greater than the variation in these pooled reductions. Thus the singular-value distribution appears to be a stronger predictor of relative work reduction than matrix size over the tested range.

8.5. Candidate screening

Two inexpensive rejection conditions assess the feasibility of one of the two rotations. They eliminate approximately 99.95% of candidates before production verification. They are designed as sufficient conditions for rejection, and no valid CAT opportunities were missed by the screens in the experiments. Of the 32,779 candidates passing the screens, only 16 were rejected by production. Of these, 8 failed the screened rotation, while 8 failed only the unexamined rotation. The narrow gap between exact and approximate rejection conditions explains their close agreement for the screened rotation when $M \geq 2\tau$. Because these few production rejections contribute negligible overhead, we leave verification of the second rotation to production rather than introduce another screening test. Screening and accepted CAT rotations together account for approximately 0.12% of total weighted work.

8.6. Numerical accuracy

In the tested matrix families, CAT maintained numerical accuracy comparable to conventional bidiagonal QR. The reported singular-value errors, left and right orthogonality errors, and reconstruction residuals remained of the same order. These results establish “Chase and Trim” as a useful addition to conventional bidiagonal QR for the tested matrix families.

Appendix A. Complete experimental results

The complete per-family numerical results are given below. All recorded baseline, CAT, deflation, screening, and accuracy quantities are retained.

Appendix A.1. Power-law singular values $\alpha = 0.5$

Table A1.1: Power-law singular values $\alpha = 0.5$: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,296.4	3,905.6	9.10%	390.8
128	17,057.4	15,331.4	10.12%	1,726.0
256	66,064.4	61,070.2	7.56%	4,994.2
512	262,866.8	244,153.2	7.12%	18,713.6
1024	1,032,371.2	958,467.6	7.16%	73,903.6

Table A1.2: Power-law singular values $\alpha = 0.5$: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,542,567.2	3,276,458.6	7.51%	266,108.6
128	27,151,463.4	24,625,831.2	9.30%	2,525,632.2
256	206,612,821.4	191,871,660.8	7.13%	14,741,160.6
512	1,629,600,902.2	1,516,952,370.4	6.91%	112,648,531.8
1024	12,742,869,197.6	11,843,637,528.2	7.06%	899,231,669.4

Table A1.3: Power-law singular values $\alpha = 0.5$: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	303	212	96	31.17%	0.0528	0.0524	0.0478
128	620	424	201	32.16%	0.0304	0.0325	0.0290
256	1253	848	410	32.59%	0.0193	0.0218	0.0175
512	2496	1740	770	30.68%	0.0177	0.0180	0.0113
1024	4977	3507	1509	30.08%	0.0179	0.0156	0.0086

Table A1.4: Power-law singular values $\alpha = 0.5$: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	34,590.8	1.06%	16,243.2	0.50%
128	134,413.6	0.55%	64,884.0	0.26%
256	529,039.2	0.28%	258,304.4	0.13%
512	2,064,734.8	0.14%	958,204.0	0.06%
1024	7,945,444.4	0.07%	3,732,094.8	0.03%

Table A1.5: Power-law singular values $\alpha = 0.5$: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	20564	6061	29.47%	14407	70.06%
128	81246	29573	36.40%	51472	63.35%
256	324056	135654	41.86%	187992	58.01%
512	1285052	620897	48.32%	663385	51.62%
1024	5009448	2678905	53.48%	2329034	46.49%

Table A1.6: Power-law singular values $\alpha = 0.5$: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	96	0
128	0	201	0
256	0	410	0
512	0	770	0
1024	0	1509	0

Table A1.7: Power-law singular values $\alpha = 0.5$: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	3.192×10^{-15}	3.442×10^{-15}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	3.192×10^{-15}	3.442×10^{-15}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	2.198×10^{-14}	2.442×10^{-14}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.314×10^{-13}	2.218×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.174×10^{-13}	2.099×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	1.027×10^{-14}	1.009×10^{-14}

Appendix A.2. Power-law singular values $\alpha = 1.0$

Table A2.1: Power-law singular values $\alpha = 1.0$: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,784.4	3,362.6	11.15%	421.8
128	15,035.6	13,303.6	11.52%	1,732.0
256	58,552.8	52,162.8	10.91%	6,390.0
512	232,961.6	208,625.8	10.45%	24,335.8
1024	909,895.6	818,932.8	10.00%	90,962.8

Table A2.2: Power-law singular values $\alpha = 1.0$: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,120,586.2	2,828,247.6	9.37%	292,338.6
128	23,934,503.8	21,400,186.2	10.59%	2,534,317.6
256	183,131,778.8	164,007,564.0	10.44%	19,124,214.8
512	1,444,247,854.4	1,296,657,031.8	10.22%	147,590,822.6
1024	11,231,265,509.0	10,120,868,709.8	9.89%	1,110,396,799.2

Table A2.3: Power-law singular values $\alpha = 1.0$: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	305	205	105	33.87%	0.0406	0.0583	0.0486
128	621	400	224	35.90%	0.0312	0.0350	0.0297
256	1234	798	452	36.16%	0.0211	0.0258	0.0177
512	2452	1630	854	34.38%	0.0243	0.0240	0.0143
1024	4916	3340	1619	32.65%	0.0287	0.0224	0.0131

Table A2.4: Power-law singular values $\alpha = 1.0$: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	31,611.0	1.12%	17,766.0	0.63%
128	123,712.4	0.58%	72,309.6	0.34%
256	476,207.4	0.29%	284,768.0	0.17%
512	1,833,240.2	0.14%	1,062,743.6	0.08%
1024	6,962,566.4	0.07%	4,004,170.8	0.04%

Table A2.5: Power-law singular values $\alpha = 1.0$: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	18619	4758	25.55%	13756	73.88%
128	74002	23728	32.06%	50050	67.63%
256	289209	110922	38.35%	177835	61.49%
512	1134008	519974	45.85%	613179	54.07%
1024	4378306	2295961	52.44%	2080726	47.52%

Table A2.6: Power-law singular values $\alpha = 1.0$: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	105	0
128	0	224	0
256	0	452	0
512	0	855	1
1024	0	1619	0

Table A2.7: Power-law singular values $\alpha = 1.0$: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	3.442×10^{-15}	3.553×10^{-15}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	3.442×10^{-15}	3.553×10^{-15}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	4.940×10^{-14}	5.262×10^{-14}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.283×10^{-13}	2.073×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.032×10^{-13}	1.938×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	1.015×10^{-14}	9.244×10^{-15}

Appendix A.3. Power-law singular values $\alpha = 2.0$

Table A3.1: Power-law singular values $\alpha = 2.0$: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,191.0	2,697.2	15.47%	493.8
128	13,213.0	10,983.0	16.88%	2,230.0
256	50,576.8	42,431.6	16.10%	8,145.2
512	197,351.6	167,936.4	14.90%	29,415.2
1024	783,066.0	663,316.6	15.29%	119,749.4

Table A3.2: Power-law singular values $\alpha = 2.0$: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	2,631,448.0	2,279,669.8	13.37%	351,778.2
128	21,017,581.6	17,702,521.8	15.77%	3,315,059.8
256	158,168,143.6	133,563,799.6	15.56%	24,604,344.0
512	1,223,435,209.2	1,044,237,140.6	14.65%	179,198,068.6
1024	9,665,657,311.8	8,199,507,488.4	15.17%	1,466,149,823.4

Table A3.3: Power-law singular values $\alpha = 2.0$: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	306	192	116	37.66%	0.0469	0.0594	0.0494
128	608	355	261	42.37%	0.0366	0.0422	0.0314
256	1200	720	510	41.46%	0.0320	0.0325	0.0209
512	2402	1534	903	37.05%	0.0374	0.0364	0.0200
1024	4789	3108	1758	36.13%	0.0449	0.0360	0.0215

Table A3.4: Power-law singular values $\alpha = 2.0$: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	28,185.0	1.24%	19,628.0	0.86%
128	111,824.4	0.63%	84,256.4	0.48%
256	420,579.8	0.31%	321,316.0	0.24%
512	1,572,962.8	0.15%	1,123,738.4	0.11%
1024	5,905,996.0	0.07%	4,347,983.2	0.05%

Table A3.5: Power-law singular values $\alpha = 2.0$: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	16379	3243	19.80%	13020	79.49%
128	66364	19077	28.75%	47026	70.86%
256	253202	87998	34.75%	164693	65.04%
512	963832	404837	42.00%	558092	57.90%
1024	3674693	1771167	48.20%	1901767	51.75%

Table A3.6: Power-law singular values $\alpha = 2.0$: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	116	0
128	0	261	0
256	0	511	1
512	0	903	0
1024	0	1759	1

Table A3.7: Power-law singular values $\alpha = 2.0$: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	4.441×10^{-15}	4.441×10^{-15}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	4.441×10^{-15}	4.441×10^{-15}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	1.544×10^{-11}	1.545×10^{-11}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.081×10^{-13}	1.887×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	1.892×10^{-13}	1.743×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	1.024×10^{-14}	1.028×10^{-14}

Appendix A.4. Power-law singular values $\alpha = 3.0$

Table A4.1: Power-law singular values $\alpha = 3.0$: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,422.4	2,549.6	25.50%	872.8
128	12,729.6	9,721.8	23.63%	3,007.8
256	49,395.2	37,723.2	23.63%	11,672.0
512	193,378.8	146,836.6	24.07%	46,542.2
1024	752,699.4	582,151.4	22.66%	170,548.0

Table A4.2: Power-law singular values $\alpha = 3.0$: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	2,807,440.4	2,156,036.0	23.20%	651,404.4
128	20,226,555.0	15,681,929.6	22.47%	4,544,625.4
256	154,384,224.8	118,765,406.4	23.07%	35,618,818.4
512	1,198,527,804.0	913,275,691.6	23.80%	285,252,112.4
1024	9,289,919,409.2	7,197,245,398.2	22.53%	2,092,674,011.0

Table A4.3: Power-law singular values $\alpha = 3.0$: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	299	167	134	44.52%	0.0527	0.0624	0.0499
128	595	341	263	43.54%	0.0458	0.0469	0.0338
256	1172	705	503	41.64%	0.0433	0.0433	0.0262
512	2327	1425	965	40.38%	0.0504	0.0483	0.0302
1024	4658	2862	1937	40.36%	0.0556	0.0500	0.0298

Table A4.4: Power-law singular values $\alpha = 3.0$: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	28,781.2	1.33%	22,676.4	1.05%
128	106,624.6	0.68%	84,906.8	0.54%
256	398,120.6	0.34%	316,914.8	0.27%
512	1,438,623.0	0.16%	1,200,910.0	0.13%
1024	5,374,753.6	0.07%	4,790,712.8	0.07%

Table A4.5: Power-law singular values $\alpha = 3.0$: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	16712	3251	19.45%	13327	79.75%
128	62844	16275	25.90%	46305	73.68%
256	238745	79051	33.11%	159191	66.68%
512	877003	349956	39.90%	526082	59.99%
1024	3321998	1512261	45.52%	1807796	54.42%

Table A4.6: Power-law singular values $\alpha = 3.0$: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	134	0
128	0	264	1
256	0	503	0
512	0	965	0
1024	0	1941	4

Table A4.7: Power-law singular values $\alpha = 3.0$: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	2.665×10^{-15}	2.665×10^{-15}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	2.665×10^{-15}	2.665×10^{-15}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	2.542×10^{-08}	2.542×10^{-08}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	1.964×10^{-13}	1.751×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	1.855×10^{-13}	1.658×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	7.850×10^{-15}	1.041×10^{-14}

Appendix A.5. IID rectangular noise

Table A5.1: IID rectangular noise: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,782.0	4,515.0	5.58%	267.0
128	19,368.6	18,698.4	3.46%	670.2
256	74,587.2	70,588.2	5.36%	3,999.0
512	295,362.8	281,240.0	4.78%	14,122.8
1024	1,136,655.4	1,100,020.6	3.22%	36,634.8

Table A5.2: IID rectangular noise: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,943,517.8	3,774,588.4	4.28%	168,929.4
128	30,830,969.6	29,963,895.2	2.81%	867,074.4
256	233,268,083.0	221,539,047.8	5.03%	11,729,035.2
512	1,831,031,181.2	1,746,596,573.2	4.61%	84,434,608.0
1024	14,029,945,321.8	13,589,874,867.8	3.14%	440,070,454.0

Table A5.3: IID rectangular noise: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	304	248	59	19.22%	0.0527	0.0525	0.0516
128	630	527	101	16.08%	0.0273	0.0321	0.0312
256	1256	1018	242	19.21%	0.0178	0.0194	0.0182
512	2527	2022	510	20.14%	0.0117	0.0119	0.0107
1024	5060	4050	1033	20.32%	0.0095	0.0076	0.0068

Table A5.4: IID rectangular noise: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	38,096.6	1.01%	9,983.2	0.26%
128	156,456.2	0.52%	32,602.8	0.11%
256	578,425.4	0.26%	152,461.6	0.07%
512	2,282,457.6	0.13%	634,648.4	0.04%
1024	8,804,638.0	0.06%	2,554,822.8	0.02%

Table A5.5: IID rectangular noise: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	22635	6616	29.23%	15960	70.51%
128	93619	30145	32.20%	63373	67.69%
256	353353	144025	40.76%	209086	59.17%
512	1407452	627390	44.58%	779552	55.39%
1024	5505138	2761526	50.16%	2742579	49.82%

Table A5.6: IID rectangular noise: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	59	0
128	0	101	0
256	0	242	0
512	0	510	0
1024	0	1033	0

Table A5.7: IID rectangular noise: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	7.390×10^{-13}	7.461×10^{-13}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	1.697×10^{-14}	1.714×10^{-14}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	1.697×10^{-14}	1.714×10^{-14}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.312×10^{-13}	2.255×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.287×10^{-13}	2.267×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	1.015×10^{-14}	9.999×10^{-15}

Appendix A.6. Uniform singular values [1.0, 10.0]

Table A6.1: Uniform singular values [1.0, 10.0]: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,511.2	4,261.2	5.54%	250.0
128	17,604.6	16,630.8	5.53%	973.8
256	68,205.6	64,736.8	5.09%	3,468.8
512	261,307.0	245,595.6	6.01%	15,711.4
1024	1,029,395.0	961,766.4	6.57%	67,628.6

Table A6.2: Uniform singular values [1.0, 10.0]: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,720,966.6	3,563,925.2	4.22%	157,041.4
128	28,028,284.8	26,667,683.0	4.85%	1,360,601.8
256	213,327,436.8	203,215,869.0	4.74%	10,111,567.8
512	1,619,992,940.6	1,525,430,433.8	5.84%	94,562,506.8
1024	12,706,443,468.8	11,883,141,602.0	6.48%	823,301,866.8

Table A6.3: Uniform singular values [1.0, 10.0]: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	300	245	58	19.14%	0.0710	0.0568	0.0569
128	611	495	124	20.03%	0.0455	0.0375	0.0370
256	1234	986	262	20.99%	0.0256	0.0222	0.0213
512	2461	1933	555	22.31%	0.0201	0.0179	0.0156
1024	4878	3743	1253	25.08%	0.0138	0.0162	0.0131

Table A6.4: Uniform singular values [1.0, 10.0]: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	35,653.2	1.00%	9,815.6	0.28%
128	137,534.2	0.52%	40,030.4	0.15%
256	526,184.2	0.26%	165,066.4	0.08%
512	1,955,428.8	0.13%	690,657.6	0.05%
1024	7,564,589.4	0.06%	3,098,955.2	0.03%

Table A6.5: Uniform singular values [1.0, 10.0]: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	21378	7068	33.06%	14252	66.67%
128	83313	31073	37.30%	52116	62.55%
256	323991	142499	43.98%	181230	55.94%
512	1228705	640630	52.14%	587519	47.82%
1024	4810469	2735637	56.87%	2073579	43.11%

Table A6.6: Uniform singular values [1.0, 10.0]: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	58	0
128	0	124	0
256	0	262	0
512	0	556	1
1024	0	1253	0

Table A6.7: Uniform singular values [1.0, 10.0]: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	1.155×10^{-13}	9.237×10^{-14}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	1.179×10^{-14}	1.239×10^{-14}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	1.179×10^{-14}	1.239×10^{-14}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.214×10^{-13}	2.164×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.155×10^{-13}	2.111×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	9.816×10^{-15}	9.605×10^{-15}

Appendix A.7. Uniform singular values [1.0, 100.0]

Table A7.1: Uniform singular values [1.0, 100.0]: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,156.2	3,921.0	5.66%	235.2
128	16,147.8	15,403.8	4.61%	744.0
256	64,118.2	60,606.4	5.48%	3,511.8
512	246,003.0	233,866.4	4.93%	12,136.6
1024	947,916.0	891,386.2	5.96%	56,529.8

Table A7.2: Uniform singular values [1.0, 100.0]: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,428,439.4	3,281,139.0	4.30%	147,300.4
128	25,708,222.8	24,695,260.2	3.94%	1,012,962.6
256	200,535,690.4	190,223,900.2	5.14%	10,311,790.2
512	1,525,098,762.0	1,452,491,551.0	4.76%	72,607,211.0
1024	11,700,486,492.2	11,012,637,470.6	5.88%	687,849,021.6

Table A7.3: Uniform singular values [1.0, 100.0]: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	304	244	61	20.00%	0.0610	0.0570	0.0572
128	617	511	109	17.58%	0.0521	0.0370	0.0370
256	1247	1018	228	18.30%	0.0308	0.0226	0.0216
512	2480	1996	501	20.06%	0.0259	0.0203	0.0174
1024	4989	4023	990	19.75%	0.0178	0.0153	0.0124

Table A7.4: Uniform singular values [1.0, 100.0]: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	32,741.8	1.00%	10,322.4	0.31%
128	125,917.0	0.51%	35,188.4	0.14%
256	483,419.2	0.25%	143,647.2	0.08%
512	1,828,679.6	0.13%	623,462.0	0.04%
1024	6,746,451.6	0.06%	2,448,500.0	0.02%

Table A7.5: Uniform singular values [1.0, 100.0]: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	19677	6692	34.01%	12924	65.68%
128	77161	32432	42.03%	44620	57.83%
256	303330	156437	51.57%	146665	48.35%
512	1170065	693632	59.28%	475931	40.68%
1024	4458332	3196365	71.69%	1260977	28.28%

Table A7.6: Uniform singular values [1.0, 100.0]: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	61	0
128	0	109	0
256	0	228	0
512	0	502	1
1024	0	990	0

Table A7.7: Uniform singular values [1.0, 100.0]: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	1.023×10^{-12}	7.674×10^{-13}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	1.683×10^{-14}	1.818×10^{-14}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	1.683×10^{-14}	1.818×10^{-14}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.132×10^{-13}	2.046×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.116×10^{-13}	2.043×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	9.659×10^{-15}	9.419×10^{-15}

Appendix A.8. Uniform singular values [1.0, 1000.0]

Table A8.1: Uniform singular values [1.0, 1000.0]: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,051.6	3,819.2	5.74%	232.4
128	16,038.2	15,175.2	5.38%	863.0
256	61,894.2	58,920.2	4.80%	2,974.0
512	244,539.2	230,531.0	5.73%	14,008.2
1024	929,653.0	876,800.0	5.69%	52,853.0

Table A8.2: Uniform singular values [1.0, 1000.0]: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,342,443.8	3,194,859.0	4.42%	147,584.8
128	25,533,440.4	24,328,268.4	4.72%	1,205,172.0
256	193,576,817.4	184,915,845.6	4.47%	8,660,971.8
512	1,516,009,464.4	1,431,719,336.2	5.56%	84,290,128.2
1024	11,475,007,701.4	10,832,198,197.8	5.60%	642,809,503.6

Table A8.3: Uniform singular values [1.0, 1000.0]: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	302	250	55	18.03%	0.0676	0.0554	0.0563
128	616	513	107	17.26%	0.0515	0.0346	0.0360
256	1256	1042	209	16.71%	0.0304	0.0215	0.0208
512	2498	2033	468	18.71%	0.0270	0.0196	0.0170
1024	5010	4095	923	18.39%	0.0174	0.0152	0.0117

Table A8.4: Uniform singular values [1.0, 1000.0]: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	31,831.4	1.00%	9,307.2	0.29%
128	123,877.4	0.51%	34,542.4	0.14%
256	465,377.4	0.25%	131,676.4	0.07%
512	1,796,168.2	0.13%	582,395.6	0.04%
1024	6,587,228.0	0.06%	2,282,794.4	0.02%

Table A8.5: Uniform singular values [1.0, 1000.0]: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	19161	6646	34.69%	12460	65.03%
128	76011	32356	42.57%	43548	57.29%
256	294863	163440	55.43%	131214	44.50%
512	1153360	699738	60.67%	453153	39.29%
1024	4385318	3265861	74.47%	1118534	25.51%

Table A8.6: Uniform singular values [1.0, 1000.0]: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	55	0
128	0	107	0
256	0	209	0
512	0	469	1
1024	0	923	0

Table A8.7: Uniform singular values [1.0, 1000.0]: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	1.205×10^{-11}	8.754×10^{-12}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	1.265×10^{-13}	1.066×10^{-13}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	1.265×10^{-13}	1.066×10^{-13}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.058×10^{-13}	1.995×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.112×10^{-13}	2.057×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	9.264×10^{-15}	9.180×10^{-15}

Appendix A.9. Uniform singular values [1.0, 10000.0]

Table A9.1: Uniform singular values [1.0, 10000.0]: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,067.6	3,859.8	5.11%	207.8
128	15,698.4	14,985.0	4.54%	713.4
256	60,544.2	58,165.2	3.93%	2,379.0
512	237,287.0	225,440.8	4.99%	11,846.2
1024	925,946.8	873,690.6	5.64%	52,256.2

Table A9.2: Uniform singular values [1.0, 10000.0]: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,351,418.4	3,224,854.6	3.78%	126,563.8
128	24,990,746.4	24,021,493.2	3.88%	969,253.2
256	189,349,337.0	182,527,163.0	3.60%	6,822,174.0
512	1,471,029,247.4	1,400,048,596.6	4.83%	70,980,650.8
1024	11,429,261,913.0	10,793,833,002.0	5.56%	635,428,911.0

Table A9.3: Uniform singular values [1.0, 10000.0]: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	306	251	54	17.70%	0.0684	0.0552	0.0569
128	616	513	108	17.39%	0.0525	0.0343	0.0360
256	1257	1057	194	15.51%	0.0311	0.0196	0.0197
512	2508	2055	449	17.93%	0.0248	0.0181	0.0155
1024	5018	4086	939	18.69%	0.0183	0.0148	0.0120

Table A9.4: Uniform singular values [1.0, 10000.0]: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	32,022.6	0.99%	9,138.0	0.28%
128	121,667.4	0.51%	34,865.2	0.15%
256	455,649.0	0.25%	122,226.0	0.07%
512	1,726,366.4	0.12%	558,750.4	0.04%
1024	6,563,224.8	0.06%	2,322,364.8	0.02%

Table A9.5: Uniform singular values [1.0, 10000.0]: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	19363	7077	36.55%	12232	63.17%
128	75061	33606	44.77%	41347	55.08%
256	291057	170634	58.63%	120229	41.31%
512	1127840	759364	67.33%	368027	32.63%
1024	4369788	3255984	74.51%	1112865	25.47%

Table A9.6: Uniform singular values [1.0, 10000.0]: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	54	0
128	0	108	0
256	0	194	0
512	0	449	0
1024	0	939	0

Table A9.7: Uniform singular values [1.0, 10000.0]: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	1.455×10^{-10}	9.823×10^{-11}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	9.955×10^{-13}	9.955×10^{-13}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	9.955×10^{-13}	9.955×10^{-13}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.126×10^{-13}	2.078×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.118×10^{-13}	2.028×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	9.567×10^{-15}	9.514×10^{-15}

Appendix A.10. Linearly spaced singular values [0.001, 1000.0]

Table A10.1: Linearly spaced singular values [0.001, 1000.0]: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,986.0	3,796.6	4.75%	189.4
128	15,434.4	14,706.8	4.71%	727.6
256	59,863.6	56,596.4	5.46%	3,267.2
512	231,107.8	217,378.6	5.94%	13,729.2
1024	891,488.2	836,976.0	6.11%	54,512.2

Table A10.2: Linearly spaced singular values [0.001, 1000.0]: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,285,979.8	3,173,316.2	3.43%	112,663.6
128	24,566,759.4	23,575,417.6	4.04%	991,341.8
256	187,212,791.8	177,694,788.6	5.08%	9,518,003.2
512	1,432,670,591.2	1,350,352,979.8	5.75%	82,317,611.4
1024	11,003,705,089.4	10,341,964,223.4	6.01%	661,740,866.0

Table A10.3: Linearly spaced singular values [0.001, 1000.0]: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	303	242	63	20.66%	0.0738	0.0465	0.0524
128	619	490	132	21.22%	0.0369	0.0294	0.0311
256	1258	906	356	28.21%	0.0187	0.0180	0.0181
512	2529	1772	771	30.32%	0.0090	0.0110	0.0104
1024	5079	3470	1629	31.95%	0.0042	0.0067	0.0059

Table A10.4: Linearly spaced singular values [0.001, 1000.0]: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	30,456.8	0.96%	10,659.6	0.34%
128	114,347.6	0.49%	42,609.6	0.18%
256	425,940.8	0.24%	224,280.0	0.13%
512	1,596,538.0	0.12%	959,432.4	0.07%
1024	6,043,847.8	0.06%	4,028,842.8	0.04%

Table A10.5: Linearly spaced singular values [0.001, 1000.0]: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	19046	9565	50.22%	9418	49.45%
128	73666	45628	61.94%	27906	37.88%
256	283338	210169	74.18%	72813	25.70%
512	1087664	903143	83.04%	183750	16.89%
1024	4186509	3729671	89.09%	455209	10.87%

Table A10.6: Linearly spaced singular values [0.001, 1000.0]: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	63	0
128	0	132	0
256	0	356	0
512	0	771	0
1024	0	1629	0

Table A10.7: Linearly spaced singular values [0.001, 1000.0]: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	1.069×10^{-11}	9.095×10^{-12}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	9.526×10^{-14}	9.526×10^{-14}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	9.526×10^{-11}	9.526×10^{-11}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	1.974×10^{-13}	1.915×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.017×10^{-13}	1.938×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	9.196×10^{-15}	8.988×10^{-15}

Appendix A.11. Geometric singular values [0.001, 1000.0]

Table A11.1: Geometric singular values [0.001, 1000.0]: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,429.4	2,411.6	45.55%	2,017.8
128	17,356.2	9,751.4	43.82%	7,604.8
256	68,979.0	39,451.8	42.81%	29,527.2
512	267,841.6	159,205.2	40.56%	108,636.4
1024	1,027,899.4	638,008.8	37.93%	389,890.6

Table A11.2: Geometric singular values [0.001, 1000.0]: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,608,020.6	2,023,094.0	43.93%	1,584,926.6
128	27,454,947.0	15,631,411.8	43.07%	11,823,535.2
256	215,043,262.8	123,771,027.2	42.44%	91,272,235.6
512	1,657,748,274.8	988,369,693.4	40.38%	669,378,581.4
1024	12,677,353,284.6	7,880,548,355.8	37.84%	4,796,804,928.8

Table A11.3: Geometric singular values [0.001, 1000.0]: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	241	96	160	62.50%	0.2179	0.2153	0.2205
128	500	245	280	53.33%	0.2049	0.1869	0.1608
256	999	535	561	51.19%	0.1768	0.1699	0.1296
512	2102	1122	1121	49.98%	0.1504	0.1509	0.0921
1024	4359	2359	2207	48.34%	0.1241	0.1317	0.0689

Table A11.4: Geometric singular values [0.001, 1000.0]: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	22,745.6	1.12%	27,093.6	1.34%
128	88,791.4	0.57%	90,426.0	0.58%
256	354,524.6	0.29%	353,499.6	0.29%
512	1,403,306.4	0.14%	1,395,095.2	0.14%
1024	5,533,069.8	0.07%	5,458,570.0	0.07%

Table A11.5: Geometric singular values [0.001, 1000.0]: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	12860	1006	7.82%	11694	90.93%
128	50447	5033	9.98%	45134	89.47%
256	202827	26410	13.02%	175856	86.70%
512	809760	135654	16.75%	672985	83.11%
1024	3224635	678183	21.03%	2544245	78.90%

Table A11.6: Geometric singular values [0.001, 1000.0]: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	160	0
128	0	280	0
256	0	561	0
512	0	1121	0
1024	0	2207	0

Table A11.7: Geometric singular values [0.001, 1000.0]: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	7.276×10^{-12}	6.139×10^{-12}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	4.959×10^{-14}	4.960×10^{-14}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	2.857×10^{-11}	2.857×10^{-11}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.409×10^{-13}	1.789×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.397×10^{-13}	1.752×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	8.635×10^{-15}	8.427×10^{-15}

Appendix A.12. Tightly clustered singular values [1.0, 1.001]

Table A12.1: Tightly clustered singular values [1.0, 1.001]: QR sweep positions. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	4,670.8	4,416.4	5.45%	254.4
128	18,217.2	16,696.2	8.35%	1,521.0
256	69,250.8	65,125.4	5.96%	4,125.4
512	261,885.8	247,925.6	5.33%	13,960.2
1024	1,003,114.0	945,428.2	5.75%	57,685.8

Table A12.2: Tightly clustered singular values [1.0, 1.001]: Weighted arithmetic work. Average across samples.

n	Standard QR	CAT	relative reduction	absolute reduction
64	3,853,005.4	3,695,297.0	4.09%	157,708.4
128	29,007,713.2	26,783,073.2	7.67%	2,224,640.0
256	216,601,912.4	204,451,136.8	5.61%	12,150,775.6
512	1,623,605,534.2	1,540,012,443.2	5.15%	83,593,091.0
1024	12,382,072,757.8	11,681,183,764.0	5.66%	700,888,993.8

Table A12.3: Tightly clustered singular values [1.0, 1.001]: Conventional and active deflations. Counts summed across samples; mean balances pooled across all splits.

n	Passive standard	Passive CAT	Active	Active fraction	Active CAT smaller child parent	Passive CAT smaller child parent	Passive standard smaller child parent
64	302	235	65	21.67%	0.0551	0.0560	0.0553
128	601	471	147	23.79%	0.0474	0.0377	0.0385
256	1228	971	276	22.13%	0.0263	0.0221	0.0212
512	2452	1890	609	24.37%	0.0184	0.0141	0.0150
1024	4895	3788	1204	24.12%	0.0140	0.0116	0.0105

Table A12.4: Tightly clustered singular values [1.0, 1.001]: CAT-specific work. Average across samples; percentages of total CAT algorithm work.

n	CAT test work	% of CAT work	Accepted CAT work	% of CAT work
64	36,915.0	1.00%	10,999.6	0.30%
128	138,206.8	0.52%	47,455.6	0.18%
256	535,033.8	0.26%	173,886.8	0.09%
512	2,030,202.8	0.13%	757,852.8	0.05%
1024	7,511,668.4	0.06%	2,977,765.2	0.03%

Table A12.5: Tightly clustered singular values [1.0, 1.001]: CAT rejection distribution. Counts summed across samples; percentages of total candidate inspections.

n	Total candidates	Test 1 rejections	% of candidates	Test 2 rejections	% of candidates
64	22155	7410	33.45%	14680	66.26%
128	83676	31025	37.08%	52504	62.75%
256	325945	129168	39.63%	196501	60.29%
512	1240357	506137	40.81%	733610	59.15%
1024	4728612	2499731	52.86%	2227673	47.11%

Table A12.6: Tightly clustered singular values [1.0, 1.001]: CAT screening verification. Counts summed across samples.

n	Missed opportunities	Sent to production	Rejected by production
64	0	65	0
128	0	147	0
256	0	276	0
512	0	610	1
1024	0	1208	4

Table A12.7: Tightly clustered singular values [1.0, 1.001]: Numerical accuracy and singular-vector quality. Maximum across all matrix sizes and samples.

Quantity	Formula	Standard QR	CAT
Max absolute σ error	$\max_i \sigma_i - \sigma_i^* $	1.488×10^{-14}	1.488×10^{-14}
Max scaled σ error (floor = 1)	$\max_i \frac{ \sigma_i - \sigma_i^* }{\max(1, \sigma_i^*)}$	1.486×10^{-14}	1.486×10^{-14}
Max relative σ error (for $\sigma_i^* \neq 0$)	$\max_i \frac{ \sigma_i - \sigma_i^* }{ \sigma_i^* }$	1.486×10^{-14}	1.486×10^{-14}
Left singular-vector orthogonality	$\ U^T U - I\ _F$	2.576×10^{-13}	2.481×10^{-13}
Right singular-vector orthogonality	$\ V^T V - I\ _F$	2.518×10^{-13}	2.451×10^{-13}
Relative SVD residual	$\frac{\ B - U \Sigma V^T\ _F}{\ B\ _F}$	9.233×10^{-15}	9.035×10^{-15}

Statements and declarations

Funding. The authors did not receive support from any organization for the submitted work.

Competing interests. The authors have no relevant financial or non-financial interests to disclose.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process. During preparation of this work, the authors used OpenAI ChatGPT and Anthropic Claude for technical discussion and critique, consideration of methodological options and literature, drafting and language editing, conversion of source material to LaTeX, and consistency checks of the text, tables, and figures. The authors reviewed and edited the AI-assisted output as needed, verified the mathematical derivations and numerical results, and take full responsibility for the content of the article.

References

- [1] D. S. Watkins, *Fundamentals of Matrix Computations*, 3rd ed., John Wiley & Sons, 2010.
- [2] LAPACK, DBDSQR: Singular Value Decomposition of a Real Bidiagonal Matrix Using the Implicit QR Algorithm. https://www.netlib.org/lapack/lapack_routine/dbdsqr.f.
- [3] LAPACK, DLAS2: Singular Values of a Real 2×2 Triangular Matrix. https://www.netlib.org/lapack/lapack_routine/dlas2.f.
- [4] LAPACK, DLARTG: Generation of a Real Plane Rotation. https://www.netlib.org/lapack/lapack_routine/dlartg.f90.
- [5] LAPACK, DLASR: Application of a Sequence of Plane Rotations to a Real Matrix. https://www.netlib.org/lapack/lapack_routine/dlasr.f.
- [6] LAPACK, DLASV2: Singular Value Decomposition of a Real 2×2 Upper-Triangular Matrix. https://www.netlib.org/lapack/lapack_routine/dlasv2.f.
- [7] LAPACK/BLAS, DROT: Application of a Plane Rotation to Two Real Vectors. https://www.netlib.org/lapack/lapack_routine/drot.f.
- [8] LAPACK, DLATMS: Generation of Random Matrices with Specified Singular Values or Eigenvalues. https://www.netlib.org/lapack/explore-html/da/d00/dlatms_8f_source.html.
- [9] J. W. Demmel, O. A. Marques, B. N. Parlett, and C. Vömel, A Testing Infrastructure for LAPACK's Symmetric Eigensolvers, LAPACK Working Note 182 (2007). <https://www.netlib.org/lapack/lawnspdf/lawn182.pdf>.