
Approximate update methods for the singular value decomposition

Mathematics Dissertation

Part C Mathematics and Computer Science

Oliver Britton

Trinity Term 2026

University of Oxford

Contents

Notation	ii
1 Introduction	1
1.1 Approximate singular values under perturbation	1
1.2 SVD updates in optimisation	1
1.3 Background and related work	2
1.4 Contributions and outline	3
2 Preliminaries	5
2.1 The singular value decomposition	5
2.2 Weyl's inequality	6
2.3 Li-Li bound	7
2.4 Schur complement	8
3 Approximating singular values of nearly diagonal matrices	9
3.1 The diagonal approximation	9
3.1.1 Approximating singular values by diagonal entries	9
3.1.2 Diagonal Li-Li bound	10
3.1.3 Worst-pair bound via repeated Li-Li	10
3.1.4 Per-pair bound via the Schur complement	15
3.2 The windowed approximation	22
3.2.1 Approximating singular values by local SVDs	22
3.2.2 Windowed Li-Li bound	23
3.2.3 Windowed per-pair bound via local diagonalisation	27
3.3 The banded approximation	30
4 Experiments	31
4.1 The approximations in practice	31
4.2 Tightness of the bounds	35
4.2.1 Diagonal method	35
4.2.2 Windowed method	38
5 Conclusion	40

Notation

Matrices

$A \in \mathbb{R}^{n \times n}$	the original matrix
$E \in \mathbb{R}^{n \times n}$	the perturbation
$A = U\Sigma V^\top$	SVD of A , with U, V orthogonal and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$
$F = U^\top EV$	perturbation in the SVD bases of A
$M = \Sigma + F$	$A + E$ in the SVD bases of A , nearly diagonal for small E
$D = \text{diag}(M)$	the diagonal of M
$F_{\text{off}} = F - \text{diag}(F)$	off-diagonal part of F (equivalently, of M)

Indices and windows

n	matrix dimension
k	singular value index, $1 \leq k \leq n$
e_k	the k th standard basis vector of $\mathbb{R}^n$
$\mathcal{W}$	the set of contiguous window indices
$\overline{\mathcal{W}}$	the off-window indices
$ \mathcal{W} $	window size
$W = M[\mathcal{W}, \mathcal{W}]$	the submatrix of entries within $\mathcal{W}$
$M[i, \mathcal{W}]$	the submatrix of entries $\{(i, j) \mid j \in \mathcal{W}\}$
$M[\mathcal{W}, j]$	the submatrix of entries $\{(i, j) \mid i \in \mathcal{W}\}$

Norms

$\ \cdot\ $	spectral norm
$\ \cdot\ _{\text{F}}$	Frobenius norm
$\ E\ = \delta$	spectral norm of the perturbation
μ, μ_k	norm of the relevant portion of the perturbation

Singular values and gaps

$\sigma_k(\cdot)$	the k th singular value, in decreasing order
η_i	Li-Li gap, from $\tilde{\sigma}_i$ to nearest singular value of the other block
g_k	per-index gap in A , $\min_{j \neq k} \sigma_k - \sigma_j $
g_k^M	per-index gap in M , $\min_{j \neq k} m_{kk} - m_{jj} $
$\hat{g}_k$	gap floored by twice the singular value, $\min(g_k, 2\sigma_k)$
ω_k	window gap, from σ_k to the nearest σ outside $\mathcal{W}$

1 Introduction

1.1 Approximate singular values under perturbation

Let A be a matrix whose singular value decomposition $A = U\Sigma V^\top$ is already known, and let $E \in \mathbb{R}^{n \times n}$ be a perturbation of spectral norm $\delta := \|E\|$. In this dissertation, we aim to approximate the singular values $\sigma_k(A + E)$ without recomputing the SVD of $A + E$ from scratch.

The core idea is that writing $A + E$ in the SVD bases of the original matrix gives

$$M := U^\top(A + E)V = \Sigma + F, \quad F := U^\top EV.$$

Since orthogonal transformations preserve both the singular values and the spectral norm, we have that $\sigma_k(M) = \sigma_k(A + E)$ and $\|F\| = \delta$. The matrix M therefore shares the singular values of $A + E$, but is *nearly diagonal*: a diagonal Σ plus a perturbation F of small spectral norm.

This nearly-diagonal structure is very revealing, as it encodes the singular values of $A + E$ in the entries near its diagonal. Later, we will see that when the gaps in the spectrum of A are large enough, the diagonal entries of M approximate $\sigma_k(M)$ to roughly $O(\mu_k^2/\text{gap})$ accuracy, where *gap* measures the separation in the spectrum of A and μ_k is the norm of the off-diagonal entries of the k th row and column of the perturbation.

This quadratic bound inspires several techniques for approximating the singular values of M using only the entries around its diagonal. We derive sharp error bounds for these approximations relating the error to the spectrum of A and the structure of E , and study how they behave in practice.




Figure 1. *The core observation.* Writing the perturbed matrix $A + E$ in the SVD basis of A , we obtain a matrix M that has the same singular values as $A + E$ and a nearly-diagonal structure.

1.2 SVD updates in optimisation

The SVD update setting arises naturally in many real-world scenarios. In several iterative algorithms, a matrix changes by a small amount on each iteration, and the algorithm requires a function of the SVD at every step. Recomputing the new

SVD from scratch is wasteful, as the previous SVD already contains much of the information about the matrix’s properties in the current iteration.

One such algorithm is the Muon algorithm [2], which is an optimiser for training large neural networks. In the Muon algorithm, we let $W_t \in \mathbb{R}^{m \times n}$ be the weight matrix of a neural network layer at iteration t , and let $G_t \in \mathbb{R}^{m \times n}$ be the gradient matrix containing the derivatives of the loss with respect to the weights. Muon modifies these weights via the update

$$M_t = \beta M_{t-1} + (1 - \beta)G_t, \quad W_{t+1} = W_t - \lambda \text{polar}(M_t),$$

where M_t is an exponentially-weighted running average of the gradient, $\text{polar}(M_t) = U_t V_t^\top$ and U_t, V_t are the singular vectors of M_t . When $(1 - \beta)(G_t - M_{t-1})$ is small, we may view M_t as a small perturbation of M_{t-1} , and an efficient SVD update algorithm could speed up Muon by reusing the SVD from the previous iteration.

1.3 Background and related work

Updating the SVD and eigendecomposition. This line of work began with computing updates to eigenvalue problems under low-rank perturbations. Golub [7] showed that the eigenvalues of a matrix plus a rank-one perturbation are the roots of a particular rational equation, which Bunch et al. [5] adapted into a practical algorithm. Bunch and Nielsen [4] adapted these eigenvalue results to the case of updating the SVD when adding or removing a column, and Gu and Eisenstat [9] derived a stable algorithm for the standalone rank-one update before Brand [3] later gave exact update formulae for the SVD under general rank- p perturbations. Another way of calculating the SVD or eigendecomposition is by an iterative diagonalisation process like in the Jacobi algorithm [10] (for the eigendecomposition) or the Kogbetliantz algorithm [13] (for the SVD). In these algorithms, a sequence of 2×2 rotations is applied to progressively reduce the off-diagonal entries of a matrix until it is sufficiently diagonal. Recently, Kenmoe et al. [12] proposed an iterative algorithm for finding the eigenvectors of a nearly diagonal matrix, based on a fixed-point equation for eigenvalues that converges when the matrix is nearly diagonal, a setting close to ours.

Perturbation bounds and expansions. The problem of bounding changes in singular values under perturbation has a long historical precedent. Weyl [23] established the bound $|\lambda_k(A + E) - \lambda_k(A)| \leq \delta$ and Mirsky [16] then later derived the singular value analogue $|\sigma_k(A + E) - \sigma_k(A)| \leq \delta$ in any unitarily invariant norm. These bounds are tight under no additional assumptions on E but when E is block off-diagonal, tighter estimates may be obtained. Following previous work in this setting, Li and Li [14] gave bounds of order $O(\mu^2/\text{gap})$. Similar results hold for eigenvectors and singular vectors. Davis and Kahan [6] showed that $\sin \theta \leq O(\mu/\text{gap})$, where θ is the angle between the perturbed and unperturbed eigenspaces, and Wedin [22] extended this to analogous bounds for singular vectors. A complementary line of work

gives formal power-series expansions of perturbed eigenvalues and singular values in the perturbation parameter. Nakatsukasa [18] provided a unified perspective of the block-perturbation bounds and first-order eigenvalue perturbation expansions.

1.4 Contributions and outline

We will prove several bounds that explain how the perturbation and spectrum of $A + E$ affect the error in the diagonal approximation $\tilde{\sigma}_k := m_{kk}$, each at increasing levels of granularity. Weyl's inequality (Corollary 2.7) gives the unconditional linear bound,

$$|\sigma_k(A + E) - m_{kk}| \leq 2\delta. \quad (1.1)$$

When δ is small relative to the gap, the Li-Li bound (Theorem 3.1) improves this to an expression that is quadratic in the off-row and off-column norm μ_k and inversely proportional to the smallest neighbouring gap $g_k = \min_{j \neq k} |\sigma_k(M) - \sigma_j(M)|$,

$$|\sigma_k(A + E) - m_{kk}| \leq \frac{2\mu_k^2}{(g_k - 2\delta) + \sqrt{(g_k - 2\delta)^2 + 4\mu_k^2}}. \quad (1.2)$$

Then, adapting an inductive argument of Shao and Nakatsukasa [19], our worst-pair bound (Theorem 3.2) weights each entry by its own gap rather than the smallest at index k , but does this at the cost of an $O(n^2)$ prefactor and uniformity across k ,

$$|\sigma_k(A + E) - m_{kk}| \leq 4n(n-1) \max_{i \neq j} \frac{m_{ij}^2}{|m_{ii} - m_{jj}| + \sqrt{|m_{ii} - m_{jj}|^2 + 4m_{ij}^2}}. \quad (1.3)$$

Building on this, our main result (Theorem 3.4) decomposes the error into a weighted sum of per-pair contributions, each multiplied by its own gap.

$$|\sigma_k(A + E) - m_{kk}| \leq \sum_{j \neq k} \frac{m_{kj}^2 + m_{jk}^2}{|m_{kk} - m_{jj}|}. \quad (1.4)$$

We extend the Li-Li and per-pair bounds to a windowed approximation (Theorem 3.6, Theorem 3.7) where each $\tilde{\sigma}_k$ is taken from the local SVD of a contiguous submatrix $M[\mathcal{W}, \mathcal{W}]$ rather than from a single diagonal entry, and show that the resulting bound depends only on the off-window blocks coupled to the perturbation. Section 4 complements the theory with numerical experiments comparing the bounds and approximations.

We focus on singular values throughout. Bounds on the singular vectors require different arguments than those we use here, which we do not pursue. We briefly discuss a banded approximation in Section 3.3, where we conjecture but do not prove an analogue of the per-pair bound in this case.

Table 1. *Summary of our error bounds for the diagonal and windowed approximations.* $\tilde{\sigma}_k = m_{kk}$ for the diagonal bounds and $\tilde{\sigma}_k = \sigma_{k-p+1}(M[\mathcal{W}, \mathcal{W}])$ for the windowed bounds, where $\mathcal{W} = \{p, \dots, p + \ell - 1\}$ is the window. g_k is the gap from σ_k to the nearest other singular value, ω_k is the gap to the nearest singular value outside $\mathcal{W}$, and $\hat{g}_k = \min(g_k, 2\sigma_k)$. μ_k and μ are the norms of the portions of the perturbation relevant to our approximation.

Bound	Bound on $ \sigma_k(A + E) - \tilde{\sigma}_k $	Assumptions
Diagonal Li-Li Theorem 3.1	$\frac{2\mu_k^2}{(g_k - 2\delta) + \sqrt{(g_k - 2\delta)^2 + 4\mu_k^2}}$	$\delta < g_k/2$
Diagonal worst-pair Theorem 3.2	$4n(n-1) \max_{i \neq j} \frac{m_{ij}^2}{ m_{ii} - m_{jj} + \sqrt{ m_{ii} - m_{jj} ^2 + 4m_{ij}^2}}$	$m_{11} > \dots > m_{nn} > 0$
Diagonal per-pair Corollary 3.5	$\sum_{j \neq k} \frac{m_{kj}^2 + m_{jk}^2}{ m_{kk} - m_{jj} }$	$\delta < \hat{g}_k/10,$ $m_{11} > \dots > m_{nn} > 0$
Windowed Li-Li Theorem 3.6	$\frac{2\mu^2}{(\omega_k - 2\delta) + \sqrt{(\omega_k - 2\delta)^2 + 4\mu^2}}$	$\delta < \omega_k/2$
Windowed per-pair Corollary 3.8	$\frac{10}{9} \sum_{j \notin \mathcal{W}} \frac{\ M[\mathcal{W}, j]\ ^2 + \ M[j, \mathcal{W}]\ ^2}{ m_{kk} - m_{jj} }$	$\forall i \in \mathcal{W}, \delta < \hat{g}_i/22,$ $m_{11} > \dots > m_{nn} > 0$

2 Preliminaries

We work over $\mathbb{R}$. Straightforward extensions of these preliminary results also hold for $\mathbb{C}$, which we omit for brevity.

2.1 The singular value decomposition

Theorem 2.1 (Singular value decomposition). *Let $A \in \mathbb{R}^{m \times n}$. Then there exist orthogonal matrices $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ such that*

$$A = U \Sigma V^\top,$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_p) \in \mathbb{R}^{m \times n}$ with $p = \min\{m, n\}$, and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$.

See Golub and Van Loan [8, Theorem 2.4.1].

We write $\sigma_k(A)$ for the k th largest singular value of A , and u_k, v_k for the k th columns of U and V , respectively.

The Jordan-Wielandt matrix characterises the singular values of A as the non-negative eigenvalues of a symmetric matrix.

Definition 2.2 (Jordan-Wielandt matrix). Let $A \in \mathbb{R}^{n \times n}$. The *Jordan-Wielandt* matrix is the block matrix

$$H(A) := \begin{pmatrix} 0 & A \\ A^\top & 0 \end{pmatrix} \in \mathbb{R}^{2n \times 2n}.$$

Theorem 2.3 (Jordan-Wielandt). *Let $A = U \Sigma V^\top$ be the SVD of $A \in \mathbb{R}^{n \times n}$, with u_k, v_k the columns of U and V . The $2n$ vectors*

$$w_k^\pm := \frac{1}{\sqrt{2}} \begin{pmatrix} u_k \\ \pm v_k \end{pmatrix},$$

form an orthonormal eigenbasis of $H(A)$ with $H(A)w_k^+ = \sigma_k w_k^+$ and $H(A)w_k^- = -\sigma_k w_k^-$. In particular, the eigenvalues of $H(A)$ are exactly $\{\pm \sigma_k(A)\}_{k=1}^n$.

Proof. Via direct calculation, we obtain

$$H(A)w_k^+ = \frac{1}{\sqrt{2}} \begin{pmatrix} A v_k \\ A^\top u_k \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \sigma_k u_k \\ \sigma_k v_k \end{pmatrix} = \sigma_k w_k^+,$$

and similarly $H(A)w_k^- = -\sigma_k w_k^-$. Orthonormality of $\{w_k^\pm\}$ follows from the orthonormality of $\{u_k\}$ and $\{v_k\}$, and the $2n$ vectors span $\mathbb{R}^{2n}$, so they form a complete eigenbasis. $\square$

The *Jordan-Wielandt* matrix is sometimes called the *Hermitian dilation* of A , hence the notation. This construction is due first to Jordan [11], but its usage as a tool for singular value perturbation theory is due to Wielandt. Our formulation here is from Stewart and Sun [20, Theorem 4.2].

In particular, this shows that the singular values are uniquely determined by A . This lets us prove that the singular values are invariant to orthogonal transformations.

Proposition 2.4 (Orthogonal invariance of singular values). *For any $A \in \mathbb{R}^{m \times n}$ and any orthogonal matrices $P \in \mathbb{R}^{m \times m}$, $Q \in \mathbb{R}^{n \times n}$, the singular values of $PAQ^\top$ coincide with those of A .*

Proof. Let $A = U\Sigma V^\top$ be an SVD of A . Then

$$PAQ^\top = PU\Sigma(QV)^\top,$$

is an SVD of $PAQ^\top$, as PU and QV are also orthogonal. Hence by uniqueness the singular values of $PAQ^\top$ are the diagonal entries of Σ . $\square$

Taking $P = U^\top$ and $Q = V^\top$ in Proposition 2.4, applied to $A + E$, establishes that M has the same singular values as $A + E$. Since $\|U^\top EV\| = \|E\| = \delta$, we can exploit the nearly-diagonal structure of M to compute the singular values of $A + E$.

Corollary 2.5 (The core observation). *Let $A = U\Sigma V^\top$ be an SVD of $A \in \mathbb{R}^{n \times n}$ and let $E \in \mathbb{R}^{n \times n}$, and write $M := \Sigma + F$ where $F := U^\top EV$. Then*

$$\sigma_k(A + E) = \sigma_k(M),$$

for all $k = 1, \dots, n$, and $\|F\| = \|E\| = \delta$.

With this observation, we now present two matrix perturbation results we use to bound how the singular values of a matrix A change under perturbation. Weyl's inequality (Theorem 2.6) gives an unconditional bound. The Li-Li bound (Theorem 2.8) sharpens this when the perturbation has block structure and the spectrum is well-separated.

2.2 Weyl's inequality

Theorem 2.6 (Weyl's inequality for singular values). *For any $A, E \in \mathbb{R}^{m \times n}$ and all $k = 1, \dots, \min(m, n)$,*

$$|\sigma_k(A + E) - \sigma_k(A)| \leq \|E\|.$$

Corollary 2.7 (Weyl's inequality in the diagonal case). *Let $X \in \mathbb{R}^{n \times n}$ and let $D \in \mathbb{R}^{n \times n}$ be a diagonal matrix whose diagonal entries $d_1 \geq d_2 \geq \dots \geq d_n \geq 0$ are in non-increasing order. Then for each $k = 1, \dots, n$,*

$$|\sigma_k(X) - d_k| \leq \|X - D\|.$$

In particular, taking $D = \text{diag}(X)$ gives $|\sigma_k(X) - x_{kk}| \leq \|X_{\text{off}}\|$, where X_{off} is the off-diagonal part of X .

Proof. Taking $A = X$ and $E = D - X$ in Theorem 2.6 gives $|\sigma_k(X) - \sigma_k(D)| \leq \|X - D\|$. Since D is diagonal with non-increasing non-negative diagonal, $\sigma_k(D) = d_k$, and the claim follows. $\square$

2.3 Li-Li bound

When A has a sufficiently large gap in its spectrum and the perturbation is off-diagonal, the perturbation behaviour can be significantly tighter than the Weyl bound suggests. To see this, consider the two matrices

$$M_1 = \begin{pmatrix} 1 & \delta \\ \delta & 1 \end{pmatrix}, \quad M_2 = \begin{pmatrix} 2 & \delta \\ \delta & 1 \end{pmatrix},$$

which we view as perturbations of $\text{diag}(1, 1)$ and $\text{diag}(2, 1)$. Calculating singular values, we see $\sigma_1(M_1) = 1 + \delta$ but $\sigma_1(M_2) = 2 + O(\delta^2)$. The matrix M_1 saturates the Weyl bound, since $\sigma_1(M_1)$ differs from 1 by exactly δ . The matrix M_2 , on account of the gap in its spectrum, has a leading singular value differing from 2 by only $O(\delta^2)$.

The Li-Li bound gives a general result about perturbations like this for arbitrary block decompositions. This particular bound originates from Li and Li [14], where it is presented as a more general theorem about perturbed Hermitian matrices. The singular value form follows by applying Theorem 2.3. The Li-Li bound improved over previous quadratic-in- δ bounds such as those given by Stewart and Sun [20, Theorem 3.12], Mathias [15, Corollary 1], and Sun [21, Corollary 3.4], which were looser and required stronger assumptions.

Theorem 2.8 (Li-Li bound for singular values). *Let*

$$M = \begin{pmatrix} M_1 & F_1 \\ F_2 & M_2 \end{pmatrix} \quad \text{and} \quad \tilde{M} = \begin{pmatrix} M_1 & \\ & M_2 \end{pmatrix},$$

be real $(m+n) \times (k+\ell)$ matrices with blocks $M_1 \in \mathbb{R}^{m \times k}$, $M_2 \in \mathbb{R}^{n \times \ell}$, $F_1 \in \mathbb{R}^{m \times \ell}$ and $F_2 \in \mathbb{R}^{n \times k}$. Write $\sigma_1 \geq \dots \geq \sigma_{\min\{m+n, k+\ell\}}$ and $\tilde{\sigma}_1 \geq \dots \geq \tilde{\sigma}_{\min\{m+n, k+\ell\}}$ for the singular values of M and $\tilde{M}$ respectively. Define $\mu = \max\{\|F_1\|, \|F_2\|\}$, and

$$\eta_i := \begin{cases} \min_{\nu_2 \in \sigma(M_2)} |\tilde{\sigma}_i - \nu_2|, & \text{if } \tilde{\sigma}_i \in \sigma(M_1) \\ \min_{\nu_1 \in \sigma(M_1)} |\tilde{\sigma}_i - \nu_1|, & \text{if } \tilde{\sigma}_i \in \sigma(M_2) \end{cases}$$

$$\eta := \min_{1 \leq i \leq \min\{m+n, k+\ell\}} \eta_i = \min_{\nu_1 \in \sigma(M_1), \nu_2 \in \sigma(M_2)} |\nu_1 - \nu_2|.$$

Then for $i = 1, 2, \dots, \min\{m+n, k+\ell\}$, we have

$$|\sigma_i(M) - \sigma_i(\tilde{M})| \leq \frac{2\mu^2}{\eta_i + \sqrt{\eta_i^2 + 4\mu^2}} \leq \frac{2\mu^2}{\eta + \sqrt{\eta^2 + 4\mu^2}}.$$

For our purposes, we interpret this bound as quantifying how much change there is in the singular values of M by ignoring its off-diagonal entries.

2.4 Schur complement

One final result we will need is the following characterisation of the eigenvalues of a block matrix in terms of its *Schur complement* [1].

Theorem 2.9 (Schur complement characterisation of eigenvalues). *Let $H \in \mathbb{R}^{(m+n) \times (m+n)}$ be the block matrix*

$$H = \begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix},$$

where $H_{11} \in \mathbb{R}^{m \times m}$, $H_{22} \in \mathbb{R}^{n \times n}$, $H_{12} \in \mathbb{R}^{m \times n}$, $H_{21} \in \mathbb{R}^{n \times m}$. For any λ not an eigenvalue of H_{11} , λ is an eigenvalue of H if and only if the Schur complement at λ

$$S(\lambda) := (H_{22} - \lambda I) - H_{21}(H_{11} - \lambda I)^{-1}H_{12}$$

is singular.

To see why this is useful, consider the case where $H_{11} = h$, which we view as an approximation to an eigenvalue λ of H under an off-diagonal perturbation by H_{12} and H_{21} . In this case, rearranging the singularity condition and taking norms gives

$$|\lambda - h| \leq \|H_{12}\| \|(\lambda I - H_{22})^{-1}\| \|H_{21}\| \approx \mu^2 \cdot \|(\lambda I - H_{22})^{-1}\|.$$

This is suggestive: if we can relate $\|(\lambda I - H_{22})^{-1}\|$ to the gap properties of the spectrum of M , we can derive $O(\mu^2/\text{gap})$ -style bounds. In Theorem 3.4, we adapt an argument of Mathias [15] that builds on this idea.

Proof. For λ not an eigenvalue of H_{11} , the block $H_{11} - \lambda I$ is invertible and $H - \lambda I$ admits the block decomposition

$$H - \lambda I = \begin{pmatrix} I & \\ H_{21}(H_{11} - \lambda I)^{-1} & I \end{pmatrix} \begin{pmatrix} H_{11} - \lambda I & \\ & S(\lambda) \end{pmatrix} \begin{pmatrix} I & (H_{11} - \lambda I)^{-1}H_{12} \\ & I \end{pmatrix}.$$

The triangular factors have determinant 1, so $\det(H - \lambda I) = \det(H_{11} - \lambda I) \det S(\lambda)$. Since $\det(H_{11} - \lambda I) \neq 0$, $\det(H - \lambda I) = 0$ if and only if $\det S(\lambda) = 0$. $\square$

3 Approximating singular values of nearly diagonal matrices

We now derive error bounds for the diagonal (Section 3.1) and windowed (Section 3.2) approximations, and briefly discuss the banded approximation (Section 3.3). Many of our results are more general than the SVD update setting and bound the singular values of any nearly-diagonal matrix $M = \Sigma + F$ with Σ diagonal and $\|F\| = \delta$.

All of our bounds tell a similar story. For small perturbations, $|\sigma_k(M) - \tilde{\sigma}_k(M)|$ depends quadratically on the norm of the entries of F that couple the index k (or the window containing it) to the rest of the spectrum, and inversely on a *gap* quantity that measures how tightly the singular values of M are packed together.

From here onwards, we work only with square matrices $A \in \mathbb{R}^{n \times n}$. Extending to the rectangular setting seems straightforward but adds a notational overhead we have avoided here.

3.1 The diagonal approximation

3.1.1 Approximating singular values by diagonal entries

The *diagonal approximation* estimates the singular values of M as the diagonal entries of M :

$$\tilde{\sigma}_k(M) := m_{kk} \quad \text{for } k = 1, \dots, n.$$

When the diagonal of M is in non-increasing order, Weyl's inequality in the diagonal case (Corollary 2.7) gives an immediate bound

$$|\sigma_k(M) - m_{kk}| \leq \|F_{\text{off}}\| \leq 2\delta,$$

where we write F_{off} for the off-diagonal entries of F . This is linear in δ and is tight under no additional assumptions, as we observed in Section 2.3. In the next three sections, we refine this to $O(\mu_k^2/\text{gap})$, where $\mu_k \leq \delta$ is the norm of the off-diagonal entries of F in row k and column k , by incorporating hypotheses about gaps in the spectrum of M .




Figure 2. *The diagonal approximation.* We approximate $\sigma_k(M)$ by the diagonal entries of M .

3.1.2 Diagonal Li-Li bound

Theorem 3.1 (Diagonal Li-Li bound for the SVD). *Let $A = U\Sigma V^\top \in \mathbb{R}^{n \times n}$ have singular values $\sigma_1 \geq \dots \geq \sigma_n \geq 0$, and let $E \in \mathbb{R}^{n \times n}$ be a perturbation with $\delta := \|E\|$. Write $M := \Sigma + F$ where $F := U^\top EV$, and assume $m_{kk} > 0$. Define the gap*

$$g_k := \min_{j \neq k} |\sigma_k - \sigma_j|,$$

and write $\mu_k := \max \left(\sqrt{\sum_{j \neq k} m_{kj}^2}, \sqrt{\sum_{j \neq k} m_{jk}^2} \right) \leq \delta$ for the max of the off-diagonal row and column norms at index k .

If $\delta < g_k/2$, then

$$|\sigma_k(A + E) - m_{kk}| \leq \frac{2\mu_k^2}{(g_k - 2\delta) + \sqrt{(g_k - 2\delta)^2 + 4\mu_k^2}}.$$

We also have the unconditional looser bound

$$|\sigma_k(A + E) - m_{kk}| \leq \frac{4\delta^2}{g_k}$$

valid for all $g_k > 0$.

We defer the proof of this theorem to Section 3.2.2, where we interpret it as the special case of a 1×1 window. Note that g_k is defined in terms of the spectrum of A rather than the spectra of the submatrices of M , and so differs from the gap η_k in the original Li-Li result (Theorem 2.8).

3.1.3 Worst-pair bound via repeated Li-Li

The Li-Li bound of the previous section pairs the off-diagonal k th row and column of the perturbation with the worst gap g_k . In this section, we show that the error can be expressed by a Li-Li-like quantity that pairs each off-diagonal entry of the perturbation with its own gap quantity. This costs an $O(n^2)$ dimension dependence and a uniform bound for all k , both of which we resolve in Theorem 3.4.

The eigenvalue analogue of this bound for symmetric matrices is due to Shao and Nakatsukasa [19]. The inductive Li-Li argument in our proof follows theirs, with the eigenvalue Li-Li bound substituted for its singular value counterpart.

Theorem 3.2 (Diagonal worst-pair bound). *Let $M \in \mathbb{R}^{n \times n}$ be a real matrix with positive diagonal entries $m_{11} \geq \dots \geq m_{nn} > 0$. Then for each $k = 1, \dots, n$,*

$$|\sigma_k(M) - m_{kk}| \leq 4n(n-1) \max_{i \neq j} \frac{m_{ij}^2}{|m_{ii} - m_{jj}| + \sqrt{|m_{ii} - m_{jj}|^2 + 4m_{ij}^2}}.$$




Figure 3. *The diagonal worst-pair bound.* One interpretation of Theorem 3.2 is that we score each pair (i, j) with $i \neq j$ by the quantity on the right-hand side of the bound. The error in the diagonal approximation across all indices is bounded by $4n(n-1)$ times the worst-pair score.

Corollary 3.3 (Diagonal worst-pair bound for the SVD). *Let $A = U\Sigma V^\top \in \mathbb{R}^{n \times n}$ have singular values $\sigma_1 \geq \dots \geq \sigma_n \geq 0$ and let $E \in \mathbb{R}^{n \times n}$, and write $M := \Sigma + F$ where $F := U^\top E V$. Assume $m_{11} > \dots > m_{nn} > 0$. Then the diagonal approximation error satisfies*

$$|\sigma_k(A + E) - m_{kk}| \leq 4n(n-1) \max_{i \neq j} \frac{m_{ij}^2}{|m_{ii} - m_{jj}| + \sqrt{|m_{ii} - m_{jj}|^2 + 4m_{ij}^2}}.$$

Unlike Li-Li, this bound pairs each off-diagonal entry m_{ij} with its own gap $|m_{ii} - m_{jj}|$ rather than the smallest gap $\min_j |m_{kk} - m_{jj}|$, and for this reason there are examples where this bound is tight but Li-Li is loose.

Consider the specific example

$$M = \begin{pmatrix} 1 + 10^{-4} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0.5 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 10^{-6} \\ 0 & 0 & 0 \\ 10^{-6} & 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 + 10^{-4} & 0 & 10^{-6} \\ 0 & 1 & 0 \\ 10^{-6} & 0 & 0.5 \end{pmatrix}.$$

By direct calculation, we see that $\sigma_1(M) - m_{11} \approx 10^{-12}$. The Li-Li bound estimates the error in σ_1 via $\mu_1 = 10^{-6}$ and gap $g_1 = 10^{-4}$. This gives the estimate $\approx \mu_1^2/g_1 = 10^{-8}$. In contrast, the terms in the worst-pair bound corresponding to the small gap are zero, and instead the maximum is attained at the entries $(1, 3)$ and $(3, 1)$, which give a right-hand side of approximately 2×10^{-11} .

The price we pay for this is the $O(n^2)$ prefactor and the fact that the bound is uniform across k . For Gaussian random perturbations, entries of F have a typical size of $\delta/\sqrt{n}$, so the maximum pair-ratio is of order $O(\delta^2/(n \cdot \text{gap}))$, and the $O(n^2)$ prefactor reduces this to $O(n\delta^2/\text{gap})$, which is looser than Li-Li by a factor of n .

Proof of Theorem 3.2. Set

$$\epsilon := \max_{i \neq j} \frac{2m_{ij}^2}{|m_{ii} - m_{jj}| + \sqrt{|m_{ii} - m_{jj}|^2 + 4m_{ij}^2}}.$$

The stated bound is $4n(n-1) \cdot \epsilon/2 = 2n(n-1)\epsilon$, which we prove by induction on n .

Base case ($n = 2$). We write

$$M = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix}, \quad \tilde{M} = \begin{pmatrix} m_{11} & 0 \\ 0 & m_{22} \end{pmatrix}.$$

We apply the original Li-Li bound (Theorem 2.8) to this matrix viewed as a 2×2 block decomposition to obtain

$$|\sigma_k(M) - m_{kk}| \leq \frac{2\mu^2}{\eta + \sqrt{\eta^2 + 4\mu^2}},$$

where $\eta = |m_{11} - m_{22}|$ and $\mu = \max\{|m_{12}|, |m_{21}|\}$. Since the Li-Li bound is monotonically increasing in μ^2 , this bound equals ϵ . Hence $|\sigma_k(M) - m_{kk}| \leq \epsilon \leq 2n(n-1)\epsilon$, and so the bound is established in the 2×2 case.

Inductive step. Assume the bound holds for matrices of size $(n-1) \times (n-1)$ and below, and suppose M is of size $n \times n$.

From the definition of ϵ , for each $i \neq j$ we have

$$\frac{2m_{ij}^2}{|m_{ii} - m_{jj}| + \sqrt{|m_{ii} - m_{jj}|^2 + 4m_{ij}^2}} \leq \epsilon.$$

Rearranging, we have

$$2m_{ij}^2 - \epsilon|m_{ii} - m_{jj}| \leq \epsilon\sqrt{|m_{ii} - m_{jj}|^2 + 4m_{ij}^2}.$$

This implies

$$m_{ij}^2 \leq \epsilon|m_{ii} - m_{jj}| + \epsilon^2 \quad \text{for all } i \neq j. \quad (\dagger)$$

To see this, note that if the left-hand side is negative, we have the (rather loose) bound $m_{ij}^2 \leq \frac{\epsilon}{2}|m_{ii} - m_{jj}| \leq \epsilon|m_{ii} - m_{jj}| + \epsilon^2$. If the left-hand side is positive, we may square both sides and rearrange to obtain $(\dagger)$.

Set $D = \text{diag}(M)$ and define the *spread* $\gamma := m_{11} - m_{nn}$. Summing $(\dagger)$ over $i \neq j$,

we have

$$\begin{aligned}
\|M - D\|_F^2 &= \sum_{i \neq j} m_{ij}^2 \\
&\leq \epsilon \left(\sum_{i \neq j} |m_{ii} - m_{jj}| \right) + n(n-1)\epsilon^2 \\
&= 2\epsilon \left(\sum_{i < j} |m_{ii} - m_{jj}| \right) + n(n-1)\epsilon^2 \\
&= 2\epsilon \left(\sum_{l=1}^{\lfloor n/2 \rfloor} (n-2l+1) (m_{ll} - m_{n-l+1, n-l+1}) \right) + n(n-1)\epsilon^2 \quad (\star) \\
&\leq 2\epsilon\gamma \left(\sum_{l=1}^{\lfloor n/2 \rfloor} (n-2l+1) \right) + n(n-1)\epsilon^2 \\
&\leq \frac{n^2}{2}\epsilon\gamma + n(n-1)\epsilon^2 \quad (\dagger)
\end{aligned}$$

where $(\star)$ regroups the sum by noting that m_{ll} appears in the sum positively in $n-l$ pairs (l, j) with $j > l$ and negatively in the $l-1$ pairs (i, l) with $i < l$, giving a coefficient $(n-2l+1)$.

Case 1: $\gamma \leq 8n(n-2)\epsilon$ (*small spread*). Since $D = \text{diag}(M)$ is diagonal and has non-decreasing positive entries, $\sigma_k(D) = m_{kk}$. Then

$$\begin{aligned}
|\sigma_k(M) - m_{kk}|^2 &\leq \|M - D\|^2 && (\text{Corollary 2.7}) \\
&\leq \|M - D\|_F^2 && (\|\cdot\| \leq \|\cdot\|_F) \\
&\leq \frac{n^2}{2}\epsilon\gamma + n(n-1)\epsilon^2. && (\text{via } (\dagger))
\end{aligned}$$

We may then substitute this into the case assumption that $\gamma \leq 8n(n-2)\epsilon$ to obtain

$$|\sigma_k(M) - m_{kk}|^2 \leq (4n^3(n-2) + n(n-1))\epsilon^2 \leq 4n^2(n-1)^2\epsilon^2 = (2n(n-1)\epsilon)^2,$$

where the second line uses $4(n-1)^2 - 4n(n-2) = 4$, so $4n^2(n-1)^2 - 4n^3(n-2) = 4n^2 \geq n(n-1)$ for $n \geq 2$, after which taking square roots yields the claim for n .

Case 2: $\gamma > 8n(n-2)\epsilon$ (*large spread*). Define $\hat{\gamma} := \gamma - 4(n-1)(n-2)\epsilon$. The case assumption gives $\epsilon < \gamma/(8n(n-2))$, so

$$\hat{\gamma} = \gamma - 4(n-1)(n-2)\epsilon > \gamma \left(1 - \frac{(n-1)(n-2)}{2n(n-2)} \right) = \frac{n+1}{2n}\gamma > \frac{\gamma}{2}.$$

Similarly, $\epsilon < \gamma/(8n(n-2)) \leq \hat{\gamma}/(4(n-2))$ implies $\gamma + \epsilon \leq 2\hat{\gamma}$.

We now split on whether m_{kk} is close in magnitude to the top or bottom of the diagonal. At least one of $m_{kk} - m_{nn} \geq \gamma/2$ or $m_{11} - m_{kk} \geq \gamma/2$ must hold, since otherwise summing both inequalities gives $\gamma = m_{11} - m_{nn} < \gamma$, a contradiction.

Subcase 2a: $m_{kk} - m_{nn} \geq \gamma/2$. We partition M to isolate the bottom row and column,

$$M = \begin{pmatrix} \hat{M} & F_1 \\ F_2 & m_{nn} \end{pmatrix}, \quad \tilde{M} = \begin{pmatrix} \hat{M} & \\ & m_{nn} \end{pmatrix},$$

where $\hat{M} \in \mathbb{R}^{(n-1) \times (n-1)}$, $F_1 \in \mathbb{R}^{(n-1) \times 1}$, and $F_2 \in \mathbb{R}^{1 \times (n-1)}$. The inductive hypothesis applied to $\hat{M}$ gives

$$|\sigma_k(\hat{M}) - m_{kk}| \leq 2(n-1)(n-2)\epsilon. \quad (*)$$

Bounding the off-diagonal blocks via $(\dagger)$,

$$\begin{aligned} \|F_1\|^2 &= \sum_{i=1}^{n-1} m_{in}^2 \leq (n-1)\epsilon(\gamma + \epsilon) \leq 2(n-1)\epsilon\hat{\gamma}, \\ \|F_2\|^2 &= \sum_{j=1}^{n-1} m_{nj}^2 \leq (n-1)\epsilon(\gamma + \epsilon) \leq 2(n-1)\epsilon\hat{\gamma}, \end{aligned}$$

so $\mu^2 := \max(\|F_1\|^2, \|F_2\|^2) \leq 2(n-1)\epsilon\hat{\gamma}$. Applying the Li-Li bound (Theorem 2.8) to M and $\tilde{M}$ at index k ,

$$|\sigma_k(M) - \sigma_k(\tilde{M})| \leq \frac{2\mu^2}{\eta_k + \sqrt{\eta_k^2 + 4\mu^2}},$$

where, in the notation of Theorem 2.8, η_k is the minimum distance from $\sigma_k(\tilde{M})$ to the spectrum of the block that $\sigma_k(\tilde{M})$ did not originate from. To determine η_k concretely, note that the singular values of $\tilde{M}$ are $\{\sigma_i(\hat{M})\}_{i=1}^{n-1} \cup \{m_{nn}\}$ sorted by magnitude. From $(*)$ and the subcase assumption $m_{kk} - m_{nn} \geq \gamma/2$,

$$\sigma_k(\hat{M}) - m_{nn} \geq m_{kk} - m_{nn} - 2(n-1)(n-2)\epsilon \geq \frac{\gamma}{2} - 2(n-1)(n-2)\epsilon = \frac{\hat{\gamma}}{2}.$$

In particular $\sigma_k(\hat{M}) > m_{nn}$, so m_{nn} is at a position $> k$ in the merged list and $\sigma_k(\tilde{M}) = \sigma_k(\hat{M})$, originating from the $\hat{M}$ block. The block that $\sigma_k(\tilde{M})$ did not originate from is (m_{nn}) which has one singular value $\{m_{nn}\}$, so $\eta_k = \sigma_k(\hat{M}) - m_{nn} \geq \hat{\gamma}/2$. Substituting this into the Li-Li bound,

$$\begin{aligned} |\sigma_k(M) - \sigma_k(\hat{M})| &\leq \frac{2\mu^2}{\hat{\gamma}/2 + \sqrt{\hat{\gamma}^2/4 + 4\mu^2}} \\ &= \frac{4\mu^2}{\hat{\gamma} + \sqrt{\hat{\gamma}^2 + 16\mu^2}} \\ &\leq \frac{2\mu^2}{\hat{\gamma}} \\ &\leq 4(n-1)\epsilon. \end{aligned}$$

Then via the triangle inequality

$$\begin{aligned} |\sigma_k(M) - m_{kk}| &\leq 4(n-1)\epsilon + 2(n-1)(n-2)\epsilon \\ &= 2n(n-1)\epsilon, \end{aligned}$$

which gives the claim for n in this case.

Subcase 2b: $m_{11} - m_{kk} \geq \gamma/2$. This case applies only for $k \geq 2$, since otherwise the hypothesis reads $0 \geq \gamma/2 > 0$. The argument largely mirrors Subcase 2a but we omit the complete details here for brevity. The key difference is that we now have to isolate the initial row and column rather than the bottom which means we have to work with the diagonal entries $m_{22}, \dots, m_{nn}$ of $\hat{M}$. This means that m_{kk} is now at position $k - 1$ in $\hat{M}$'s diagonal, so the relevant index in $\hat{M}$'s singular values is $\sigma_{k-1}(\hat{M})$ rather than $\sigma_k(\hat{M})$. $\square$

Proof of Corollary 3.3. Apply Theorem 3.2 to M and note $\sigma_k(M) = \sigma_k(A + E)$ by Corollary 2.5. $\square$

3.1.4 Per-pair bound via the Schur complement

The worst-pair bound of the previous section demonstrates that the error in the diagonal approximation depends on how each off-diagonal entry (i, j) pairs with its own gap corresponding to m_{ii} and m_{jj} , rather than the smallest gap at index k . However, this result has an $O(n^2)$ prefactor, and the bound is uniform across all k . In this section, we remove both at the expense of stronger assumptions.

Theorem 3.4 (Diagonal per-pair bound). *Let $M \in \mathbb{R}^{n \times n}$ have strictly decreasing positive diagonal entries $m_{11} > m_{22} > \dots > m_{nn} > 0$, and write $M = D + F_{\text{off}}$ where $D := \text{diag}(M)$ and $F_{\text{off}} := M - D$ is the off-diagonal part. Fix $k \in \{1, \dots, n\}$ and define*

$$g_k^M := \min_{j \neq k} |m_{kk} - m_{jj}|, \quad \hat{g}_k^M := \min(g_k^M, 2m_{kk}).$$

If $\|F_{\text{off}}\| < \hat{g}_k^M/4$, then

$$|\sigma_k(M) - m_{kk}| \leq \sum_{j \neq k} \frac{m_{kj}^2 + m_{jk}^2}{|m_{kk} - m_{jj}|}.$$

Corollary 3.5 (Diagonal per-pair bound for the SVD). *Let $A = U\Sigma V^\top \in \mathbb{R}^{n \times n}$ have singular values $\sigma_1 \geq \dots \geq \sigma_n \geq 0$, let $E \in \mathbb{R}^{n \times n}$ be a perturbation with $\delta := \|E\|$, and write $M := \Sigma + F$ where $F := U^\top E V$. Fix $k \in \{1, \dots, n\}$ and define*

$$g_k := \min_{j \neq k} |\sigma_k - \sigma_j|, \quad \hat{g}_k := \min(g_k, 2\sigma_k).$$

If $\delta < \hat{g}_k/10$ and $m_{11} > m_{22} > \dots > m_{nn} > 0$, then

$$|\sigma_k(A + E) - m_{kk}| \leq \sum_{j \neq k} \frac{m_{kj}^2 + m_{jk}^2}{|m_{kk} - m_{jj}|}.$$

Among the three diagonal bounds, the per-pair bound is the finest in terms of its understanding of the perturbation at the cost of a slightly stronger assumption on the size of the perturbation relative to the gap.




Figure 4. *The diagonal per-pair bound.* One interpretation of Theorem 3.4 is that we score each pair (i, j) with $i \neq j$ by the quantity in the middle panel. The error in the diagonal approximation at index k is the sum of these off-diagonal quantities in row k and column k .

We now return to the example of Section 3.1.3 at $k = 1$. The small-perturbation Li-Li bound uses $\mu_1 = 10^{-6}$ paired with the small gap $g_1 = 10^{-4}$ to estimate the error as approximately $\mu_1^2/g_1 = 10^{-8}$, whereas the per-pair bound notices that the perturbation couples σ_1 only to σ_3 rather than to its small-gap neighbour σ_2 , so the relevant weight is $|m_{11} - m_{33}| \approx 0.5$ rather than g_1 . This gives a bound of approximately 4×10^{-12} , which is three orders of magnitude tighter than Li-Li in this case.

We can give a heuristic argument to show the per-pair bound is typically sharper than both the worst-pair and Li-Li bounds for typical dense perturbations. A Gaussian random matrix of norm δ has $|m_{ij}| \approx \delta/\sqrt{n}$, and its row spectral norm is $\mu_k \approx \delta/2$, so

$$\underbrace{\sum_{j \neq k} \frac{m_{kj}^2 + m_{jk}^2}{|m_{kk} - m_{jj}|}}_{\text{per-pair}} \lesssim \frac{2\mu_k^2}{n} \sum_{j \neq k} \frac{1}{|m_{kk} - m_{jj}|} \lesssim \underbrace{\frac{\mu_k^2}{g_k}}_{\text{Li-Li}}.$$

Here, we see that Li-Li weights the row and column at index k by the worst gap g_k , whereas the per-pair bound weights each entry by its own gap, and the ratio between these two behaviours is approximately $\frac{2}{n-1} \sum_{j \neq k} g_k/|m_{kk} - m_{jj}|$, which is smaller than 1 whenever most pairwise gaps are larger than g_k . For a linear spectrum with $\sigma_k = n - k + 1$, this ratio is roughly $(\log n)/n$ and hence the per-pair bound improves over Li-Li by roughly a factor of $n/\log n$.

Proof of Theorem 3.4. The idea of this proof is to write

$$\sigma_k(M) - m_{kk} = T_k + r_k,$$

where T_k is a leading-order term and r_k is a remainder, and then show that $|T_k|, |r_k| \leq S_k/2$, where S_k is the right-hand side of the claimed bound. To do this, we work with the Jordan-Wielandt matrix $H(M)$, which has $\sigma_k(M)$ as an eigenvalue, and calculate the Schur complement. This produces an implicit equation for $\sigma_k(M) - m_{kk}$ that we expand to obtain T_k and r_k .

The Schur complement argument we use on the Jordan-Wielandt matrix is due to Mathias [15], who applied it to derive an $O(\mu^2/\text{gap})$ bound on singular values. The Mathias proof bounds the expression for $\sigma_k - m_{kk}$ by appealing to submultiplicativity of the spectral norm. Our proof instead evaluates this expression coordinate-by-coordinate to obtain a sum of per-pair contributions, where each off-diagonal entry pairs with its own gap.

Setup. By Corollary 2.7 applied to M and $D := \text{diag}(M)$ (whose diagonal entries are non-negative and non-increasing by hypothesis), we obtain the loose bound

$$|\sigma_k(M) - m_{kk}| \leq \|F_{\text{off}}\|. \quad (\dagger)$$

We also write

$$S_k := \sum_{j \neq k} \frac{m_{jk}^2 + m_{kj}^2}{|m_{kk} - m_{jj}|},$$

for the right-hand side of the claimed bound.

Schur complement of the Jordan-Wielandt matrix. Write $H(M) = H(D) + H(F_{\text{off}})$ where:

$$H(D) = \begin{pmatrix} 0 & D \\ D & 0 \end{pmatrix} \in \mathbb{R}^{2n \times 2n}, \quad H(F_{\text{off}}) = \begin{pmatrix} 0 & F_{\text{off}} \\ F_{\text{off}}^\top & 0 \end{pmatrix} \in \mathbb{R}^{2n \times 2n}.$$

By the Jordan-Wielandt theorem (Theorem 2.3) the eigenvalues of $H(D)$ are $\pm m_{11}, \dots, \pm m_{nn}$ with eigenvectors $w_j^\pm := \frac{1}{\sqrt{2}} \begin{pmatrix} e_j \\ \pm e_j \end{pmatrix}$. Let $\mathcal{V} := (w_k^+)^{\perp} \subset \mathbb{R}^{2n}$ be the $(2n - 1)$ -dimensional complement of the line spanned by w_k^+ . An orthonormal basis of $\mathcal{V}$ is given by

$$\{w_j^+ : j \neq k\} \cup \{w_j^- : 1 \leq j \leq n\}.$$

Collect these $2n - 1$ vectors as the columns of $P \in \mathbb{R}^{2n \times (2n-1)}$ and define

$$\begin{aligned} \Lambda_{\perp} &:= P^\top H(D)P && \in \mathbb{R}^{(2n-1) \times (2n-1)}, \\ W_{\perp} &:= P^\top H(F_{\text{off}})P && \in \mathbb{R}^{(2n-1) \times (2n-1)}, \\ c_k &:= P^\top H(F_{\text{off}})w_k^+ && \in \mathbb{R}^{2n-1}. \end{aligned}$$

Then $\Lambda_{\perp}$ is diagonal with eigenvalues $\{\pm m_{jj}\}_{j=1}^n \setminus \{+m_{kk}\}$. Since P has orthonormal columns, $\|W_{\perp}\| \leq \|F_{\text{off}}\|$, and in the orthonormal basis $\{w_k^+\} \cup \{\text{columns of } P\}$, the matrix $H(M)$ has the block form

$$\begin{pmatrix} m_{kk} & c_k^\top \\ c_k & \Lambda_{\perp} + W_{\perp} \end{pmatrix}.$$

To see this, note that the $(1,1)$ entry is $(w_k^+)^{\top} H(M) w_k^+ = (w_k^+)^{\top} H(D) w_k^+ + (w_k^+)^{\top} H(F_{\text{off}}) w_k^+ = m_{kk} + 0$, where the second term vanishes because F_{off} has zero diagonal. The off-diagonal blocks follow from $P^\top H(D) w_k^+ = 0$ (since w_k^+ is orthogonal to $\mathcal{V}$ and $H(D) w_k^+ = m_{kk} w_k^+$), giving $c_k = P^\top H(F_{\text{off}}) w_k^+$.

By Jordan-Wielandt, $\sigma_k(M)$ is an eigenvalue of $H(M)$. Applying the Schur complement characterisation (Theorem 2.9) to this block form at this eigenvalue gives the identity

$$\sigma_k(M) - m_{kk} = c_k^\top G(\sigma_k(M)) c_k, \quad G(\lambda) := (\lambda I - \Lambda_\perp - W_\perp)^{-1}, \quad (\star)$$

provided $\sigma_k(M)$ is not an eigenvalue of $\Lambda_\perp + W_\perp$. To see why, note that the eigenvalues λ_ℓ of $\Lambda_\perp$ come in three types, and for each we have the bounds

$$\begin{aligned} \lambda_\ell = m_{jj} \ (j \neq k) : \quad & |m_{kk} - \lambda_\ell| = |m_{kk} - m_{jj}| \geq g_k^M \geq \hat{g}_k^M, \\ \lambda_\ell = -m_{jj} \ (j \neq k) : \quad & |m_{kk} - \lambda_\ell| = m_{kk} + m_{jj} \geq |m_{kk} - m_{jj}| \geq g_k^M \geq \hat{g}_k^M, \\ \lambda_\ell = -m_{kk} : \quad & |m_{kk} - \lambda_\ell| = 2m_{kk} \geq \hat{g}_k^M. \end{aligned}$$

So $\hat{g}_k^M$ is the gap in the spectrum of the Hermitian dilation $H(D)$ at the eigenvalue $+m_{kk}$.

From the first two rows, we have the sharper bound

$$|m_{kk} - \lambda_\ell| \geq g_k^M \quad \text{for } \lambda_\ell \in \lambda(\Lambda_\perp) \setminus \{-m_{kk}\}, \quad (\dagger_1)$$

and including the third we have the uniform but weaker bound

$$|m_{kk} - \lambda_\ell| \geq \hat{g}_k^M \quad \text{for every } \lambda_\ell \in \lambda(\Lambda_\perp). \quad (\dagger_2)$$

By Weyl's inequality, for every eigenvalue λ'_ℓ of $\Lambda_\perp + W_\perp$ there is a $\lambda_\ell \in \lambda(\Lambda_\perp)$ with $|\lambda_\ell - \lambda'_\ell| \leq \|W_\perp\| \leq \|F_{\text{off}}\|$. The reverse triangle inequality with $(\dagger)$, $(\dagger_2)$, and this bound gives

$$\begin{aligned} |\sigma_k(M) - \lambda'_\ell| &\geq |m_{kk} - \lambda'_\ell| - |\sigma_k(M) - m_{kk}| && \text{(reverse triangle)} \\ &\geq |m_{kk} - \lambda_\ell| - |\lambda_\ell - \lambda'_\ell| - \|F_{\text{off}}\| && \text{(reverse triangle, } (\dagger)) \\ &\geq \hat{g}_k^M - 2\|F_{\text{off}}\| && ((\dagger_2), \text{Weyl}) \\ &> \frac{1}{2}\hat{g}_k^M && (\|F_{\text{off}}\| < \hat{g}_k^M/4). \end{aligned}$$

So the candidate eigenvalue $\sigma_k(M)$ is not an eigenvalue of $\Lambda_\perp + W_\perp$ and $(\star)$ applies.

The leading term. Define

$$T_k := c_k^\top R(m_{kk}) c_k, \quad R(\lambda) := (\lambda I - \Lambda_\perp)^{-1}.$$

We first determine a concrete expression for c_k in terms of the entries of M . Since the columns of P are the orthonormal basis $\{w_j^+ : j \neq k\} \cup \{w_j^- : 1 \leq j \leq n\}$ of $\mathcal{V}$, the components of the coordinate vector $c_k = P^\top H(F_{\text{off}}) w_k^+ \in \mathbb{R}^{2n-1}$ are computed as inner products in $\mathbb{R}^{2n}$. Explicitly, one such inner product is

$$\begin{aligned} (w_j^+)^\top H(F_{\text{off}}) w_k^+ &= \frac{1}{2} \begin{pmatrix} e_j^\top & e_j^\top \end{pmatrix} \begin{pmatrix} 0 & F_{\text{off}} \\ F_{\text{off}}^\top & 0 \end{pmatrix} \begin{pmatrix} e_k \\ e_k \end{pmatrix} \\ &= \frac{1}{2} (e_j^\top F_{\text{off}} e_k + e_j^\top F_{\text{off}}^\top e_k) \\ &= \frac{1}{2} ((F_{\text{off}})_{jk} + (F_{\text{off}})_{kj}) \\ &= \frac{1}{2} (m_{jk} + m_{kj}), \end{aligned}$$

where the last equality uses $(F_{\text{off}})_{jk} = m_{jk}$ for $j \neq k$ (and the $j = k$ case gives zero because $(F_{\text{off}})_{kk} = 0$). Similarly, we may derive all the components

$$\begin{aligned} c_{jk}^+ &:= (w_j^+)^\top H(F_{\text{off}}) w_k^+ = \frac{1}{2}(m_{jk} + m_{kj}), \\ c_{jk}^- &:= (w_j^-)^\top H(F_{\text{off}}) w_k^+ = \frac{1}{2}(m_{jk} - m_{kj}), \\ c_{kk}^- &:= 0. \end{aligned}$$

We note also that

$$(c_{jk}^+)^2 + (c_{jk}^-)^2 = \frac{1}{2} (m_{jk}^2 + m_{kj}^2). \quad (\S)$$

Since $R(m_{kk})$ is diagonal in the $\{w_j^\pm\}$ -basis with entries $1/(m_{kk} - \lambda_\ell)$ over the eigenvalues λ_ℓ of $\Lambda_\perp$,

$$\begin{aligned} |T_k| &= \left| \sum_{j \neq k} \frac{(c_{jk}^+)^2}{m_{kk} - m_{jj}} + \sum_{j \neq k} \frac{(c_{jk}^-)^2}{m_{kk} + m_{jj}} \right| \quad (\text{as } c_{kk}^- = 0) \\ &\leq \sum_{j \neq k} \frac{(c_{jk}^+)^2}{|m_{kk} - m_{jj}|} + \sum_{j \neq k} \frac{(c_{jk}^-)^2}{m_{kk} + m_{jj}} \quad (\text{triangle inequality})^1 \\ &\leq \sum_{j \neq k} \frac{(c_{jk}^+)^2 + (c_{jk}^-)^2}{|m_{kk} - m_{jj}|} \quad (m_{kk} + m_{jj} \geq |m_{kk} - m_{jj}|) \\ &= \frac{1}{2} \sum_{j \neq k} \frac{m_{jk}^2 + m_{kj}^2}{|m_{kk} - m_{jj}|} \quad (\text{via } (\S)) \\ &= \frac{S_k}{2}. \end{aligned}$$

The remainder. Define $r_k := (\sigma_k(M) - m_{kk}) - T_k$. By $(\star)$ and the definition of T_k , we write

$$r_k = c_k^\top (G(\sigma_k(M)) - R(m_{kk})) c_k.$$

The resolvent identity $A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}$ applied with $A^{-1} = G(\sigma_k(M))$ and $B^{-1} = R(m_{kk})$ (so $B - A = (m_{kk} - \sigma_k(M))I + W_\perp$) yields

$$r_k = c_k^\top G(\sigma_k(M)) \left((m_{kk} - \sigma_k(M))I + W_\perp \right) R(m_{kk}) c_k.$$

Since $G(\sigma_k(M))$ is symmetric (the matrix $\Lambda_\perp + W_\perp$ is symmetric, and so is its inverse), we may rewrite $c_k^\top G(\sigma_k(M)) = (G(\sigma_k(M))c_k)^\top$. This, together with Cauchy-Schwarz, yields

$$|r_k| \leq \underbrace{\|G(\sigma_k(M))c_k\|}_{(i)} \cdot \underbrace{\|(m_{kk} - \sigma_k(M))I + W_\perp\|}_{(ii)} \cdot \underbrace{\|R(m_{kk})c_k\|}_{(iii)}.$$

We bound each factor.

¹We note that while this step yields a nice symmetric expression, it likely loosens the bound more than necessary since many of these terms may cancel. We conjecture this explains behaviour we see later in Figure 10.

(i) $\|G(\sigma_k(M))c_k\|$: For the first factor, we start by relating $G(\sigma_k(M))$ to $R(\sigma_k(M))$. The resolvent identity applied with $A^{-1} = G(\sigma_k(M))$ and $B^{-1} = R(\sigma_k(M))$ gives

$$G(\sigma_k(M)) = R(\sigma_k(M)) + R(\sigma_k(M))W_\perp G(\sigma_k(M)).$$

Applying both sides to c_k and using the triangle inequality and submultiplicativity of norms, we obtain

$$\|G(\sigma_k(M))c_k\| \leq \|R(\sigma_k(M))c_k\| + \alpha \cdot \|G(\sigma_k(M))c_k\|, \quad \alpha := \|R(\sigma_k(M))\| \|W_\perp\|.$$

We now bound α . For each eigenvalue λ_ℓ of $\Lambda_\perp$, we have

$$\begin{aligned} |\sigma_k(M) - \lambda_\ell| &\geq |m_{kk} - \lambda_\ell| - |\sigma_k(M) - m_{kk}| \quad (\text{reverse triangle inequality}) \\ &\geq \hat{g}_k^M - \|F_{\text{off}}\| \quad (\text{via } (\dagger), (\dagger_2)) \\ &> \frac{3}{4}\hat{g}_k^M. \quad (\|F_{\text{off}}\| < \hat{g}_k^M/4) \end{aligned}$$

Since $R(\sigma_k(M))$ is diagonal in the $\{w_j^\pm\}$ -basis and $\|\cdot\| = \sigma_1(\cdot)$,

$$\begin{aligned} \|R(\sigma_k(M))\| &= \frac{1}{\min_\ell |\sigma_k(M) - \lambda_\ell|} \\ &\leq \frac{4}{3\hat{g}_k^M}, \end{aligned}$$

Combined with $\|W_\perp\| \leq \|F_{\text{off}}\|$,

$$\alpha = \|R(\sigma_k(M))\| \|W_\perp\| \leq \frac{4\|F_{\text{off}}\|}{3\hat{g}_k^M} < \frac{4(\hat{g}_k^M/4)}{3\hat{g}_k^M} = \frac{1}{3}.$$

Since $\alpha < 1$, rearranging the initial inequality, we obtain

$$\|G(\sigma_k(M))c_k\| \leq \frac{\|R(\sigma_k(M))c_k\|}{1 - \alpha} \leq \frac{3}{2}\|R(\sigma_k(M))c_k\|.$$

Now that we have related $G(\sigma_k(M))$ to $R(\sigma_k(M))$, we next relate $R(\sigma_k(M))$ to $R(m_{kk})$. Both these quantities are diagonal in the $\{w_j^\pm\}$ -basis, so their componentwise ratio is

$$\frac{|R(\sigma_k(M))_{\ell\ell}|}{|R(m_{kk})_{\ell\ell}|} = \frac{|m_{kk} - \lambda_\ell|}{|\sigma_k(M) - \lambda_\ell|}.$$

We now calculate

$$\begin{aligned} |\sigma_k(M) - \lambda_\ell| &\geq |m_{kk} - \lambda_\ell| - \|F_{\text{off}}\| \quad (\text{reverse triangle, } (\dagger)) \\ &\geq |m_{kk} - \lambda_\ell| - \frac{1}{4}|m_{kk} - \lambda_\ell| \quad (\|F_{\text{off}}\| < \hat{g}_k^M/4 \leq \frac{1}{4}|m_{kk} - \lambda_\ell| \text{ via } (\dagger_2)) \\ &= \frac{3}{4}|m_{kk} - \lambda_\ell|. \end{aligned}$$

Rearranging, we see the componentwise ratio is at most $4/3$. Hence

$$\|R(\sigma_k(M))c_k\|^2 = \sum_\ell |R(\sigma_k(M))_{\ell\ell}(c_k)_\ell|^2 \leq \left(\frac{4}{3}\right)^2 \|R(m_{kk})c_k\|^2,$$

and taking square roots, $\|R(\sigma_k(M))c_k\| \leq \frac{4}{3}\|R(m_{kk})c_k\|$.

Combining these relations,

$$\|G(\sigma_k(M))c_k\| \leq \frac{3}{2}\|R(\sigma_k(M))c_k\| \leq \frac{3}{2} \cdot \frac{4}{3}\|R(m_{kk})c_k\| = 2\|R(m_{kk})c_k\|.$$

(ii) $\|(m_{kk} - \sigma_k(M))I + W_\perp\|$: For the second factor, we note

$$\begin{aligned} \|(m_{kk} - \sigma_k(M))I + W_\perp\| &\leq |m_{kk} - \sigma_k(M)| + \|W_\perp\| \\ &\leq |\sigma_k(M) - m_{kk}| + \|F_{\text{off}}\| \\ &\leq 2\|F_{\text{off}}\|, \end{aligned}$$

by the triangle inequality and $(\dagger)$.

(iii) $\|R(m_{kk})c_k\|$: For the third factor, since $R(m_{kk})$ is diagonal in the $\{w_j^\pm\}$ -basis, we obtain

$$\begin{aligned} \|R(m_{kk})c_k\|^2 &= \sum_{j \neq k} \frac{(c_{jk}^+)^2}{(m_{kk} - m_{jj})^2} + \sum_{j \neq k} \frac{(c_{jk}^-)^2}{(m_{kk} + m_{jj})^2} \\ &\leq \frac{1}{g_k^M} \left(\sum_{j \neq k} \frac{(c_{jk}^+)^2}{|m_{kk} - m_{jj}|} + \sum_{j \neq k} \frac{(c_{jk}^-)^2}{m_{kk} + m_{jj}} \right) \quad (\text{via } (\dagger_1)) \\ &\leq \frac{1}{g_k^M} \sum_{j \neq k} \frac{(c_{jk}^+)^2 + (c_{jk}^-)^2}{|m_{kk} - m_{jj}|} \quad (m_{kk} + m_{jj} \geq |m_{kk} - m_{jj}|) \\ &= \frac{S_k}{2g_k^M} \quad (\text{via } (\S)). \end{aligned}$$

Finally, combining (i), (ii) and (iii),

$$\begin{aligned} |r_k| &\leq \|G(\sigma_k(M))c_k\| \cdot \|(m_{kk} - \sigma_k(M))I + W_\perp\| \cdot \|R(m_{kk})c_k\| \\ &\leq 2\|R(m_{kk})c_k\| \cdot 2\|F_{\text{off}}\| \cdot \|R(m_{kk})c_k\| \quad (\text{by (i), (ii)}) \\ &= 4\|F_{\text{off}}\| \cdot \|R(m_{kk})c_k\|^2 \\ &\leq \frac{4\|F_{\text{off}}\| S_k}{2g_k^M} \quad (\text{by (iii)}) \\ &\leq \frac{S_k}{2}. \end{aligned}$$

where in the final inequality we use the fact that $\|F_{\text{off}}\| < g_k^M/4$ since $\hat{g}_k^M \leq g_k^M$.

By definition $\sigma_k(M) - m_{kk} = T_k + r_k$, so the triangle inequality with the bounds $|T_k| \leq S_k/2$ and $|r_k| \leq S_k/2$ just established gives

$$|\sigma_k(M) - m_{kk}| \leq |T_k| + |r_k| \leq \frac{S_k}{2} + \frac{S_k}{2} = S_k = \sum_{j \neq k} \frac{m_{kj}^2 + m_{jk}^2}{|m_{kk} - m_{jj}|},$$

which is the stated bound. $\square$

Proof of Corollary 3.5. Write $F_{\text{off}} := F - \text{diag}(F) = M - \text{diag}(M)$. Because the diagonal entry of any matrix is bounded in modulus by its spectral norm,

$$|m_{kk} - \sigma_k| = |f_{kk}| \leq \|F\| = \delta \quad \text{for each } k = 1, \dots, n, \quad (*)$$

and consequently

$$\|F_{\text{off}}\| \leq \|F\| + \|\text{diag}(F)\| \leq 2\delta.$$

Two applications of $(*)$ via the reverse triangle inequality give

$$\begin{aligned} |m_{kk} - m_{jj}| &= |(\sigma_k - \sigma_j) + (f_{kk} - f_{jj})| \\ &\geq |\sigma_k - \sigma_j| - |f_{kk} - f_{jj}| \\ &\geq |\sigma_k - \sigma_j| - 2\delta \end{aligned}$$

for each $j \neq k$, so $g_k^M \geq g_k - 2\delta$. A single application gives $m_{kk} \geq \sigma_k - \delta$. By definition of $\hat{g}_k$, the hypothesis $\delta < \hat{g}_k/10$ is equivalent to $\delta < g_k/10$ and $\delta < \sigma_k/5$. The first gives $g_k^M > 8\delta \geq 4\|F_{\text{off}}\|$. The second gives $m_{kk} > 4\delta$, so $2m_{kk} > 8\delta \geq 4\|F_{\text{off}}\|$. Hence $\hat{g}_k^M := \min(g_k^M, 2m_{kk}) > 4\|F_{\text{off}}\|$, i.e. $\|F_{\text{off}}\| < \hat{g}_k^M/4$. Theorem 3.4 applied to M at index k , together with the identity $\sigma_k(A + E) = \sigma_k(M)$ of Corollary 2.5, gives the claimed bound. $\square$

3.2 The windowed approximation

3.2.1 Approximating singular values by local SVDs




Figure 5. *The windowed approximation.* We approximate $\sigma_k(M)$ by the local SVDs of windows placed along the diagonal of M .

The *windowed approximation* generalises the diagonal approximation by considering the local SVDs of windows placed along the diagonal of M . Fix a contiguous window $\mathcal{W} = \{p, p+1, \dots, p+\ell-1\} \subset \{1, \dots, n\}$ and let $W := M[\mathcal{W}, \mathcal{W}]$ be the corresponding $\ell \times \ell$ submatrix of M . The windowed approximation estimates the singular values of M at the indices of $\mathcal{W}$ by the SVD of W :

$$\tilde{\sigma}_k(M) := \sigma_{k-p+1}(W) \quad \text{for } k \in \mathcal{W}.$$

We recover the diagonal approximation when $\ell = 1$ and the full SVD when $\ell = n$. As in Section 3.1, the error is $O(\mu^2/\text{gap})$, but now μ is the norm of the horizontally

or vertically adjacent blocks, and *gap* measures how close the singular values within $\mathcal{W}$ are to those outside $\mathcal{W}$. This means the windowed approximation can be much better than the diagonal approximation when the window spans an isolated group of very tightly clustered singular values.

3.2.2 Windowed Li-Li bound

We now state a more general version of Theorem 3.1.

Theorem 3.6 (Windowed Li-Li bound for the SVD). *Let $A = U\Sigma V^\top \in \mathbb{R}^{n \times n}$ have singular values $\sigma_1 \geq \dots \geq \sigma_n \geq 0$ and let $E \in \mathbb{R}^{n \times n}$ be a perturbation with $\delta := \|E\|$, and write $M := \Sigma + F$ where $F := U^\top E V$. Fix a contiguous window $\mathcal{W} = \{p, p+1, \dots, p+\ell-1\} \subseteq \{1, \dots, n\}$ of size $\ell := |\mathcal{W}|$, let $W := M[\mathcal{W}, \mathcal{W}]$ denote the corresponding $\ell \times \ell$ submatrix of M . For each $k \in \mathcal{W}$ write $\tilde{\sigma}_k := \sigma_{k-p+1}(W)$ for the local SVD approximation, and define the window boundary gap*

$$\omega_k := \min_{r \notin \mathcal{W}} |\sigma_k - \sigma_r|,$$

with the conventions $\sigma_0 := +\infty$ and $\sigma_{n+1} := 0$, and write

$$\mu := \max(\|M[\mathcal{W}, \overline{\mathcal{W}}]\|, \|M[\overline{\mathcal{W}}, \mathcal{W}]\|)$$

for the max of the adjacent block norms.

If $\delta < \omega_k/2$, then

$$|\sigma_k(A + E) - \tilde{\sigma}_k| \leq \frac{2\mu^2}{(\omega_k - 2\delta) + \sqrt{(\omega_k - 2\delta)^2 + 4\mu^2}}.$$

We also have the unconditional looser bound

$$|\sigma_k(A + E) - \tilde{\sigma}_k| \leq \frac{4\delta^2}{\omega_k},$$

valid whenever $\omega_k > 0$.

Compared to Theorem 3.1, the gap is now the window gap $\omega_k = \min_{r \notin \mathcal{W}} |\sigma_k - \sigma_r|$ rather than $g_k = \min_{j \neq k} |\sigma_k - \sigma_j|$, and μ is bounded by the off-window block norms rather than the full δ . To illustrate, consider the 4×4 example

$$M = \begin{pmatrix} 100 & 10^{-5} & 0 & 10^{-5} \\ 10^{-5} & 100 - 10^{-3} & 0 & 0 \\ 0 & 0 & 98 & 0 \\ 10^{-5} & 0 & 0 & 1 \end{pmatrix}, \quad \mathcal{W} = \{1, 2\},$$

which we view as a perturbation of $\text{diag}(100, 100 - 10^{-3}, 98, 1)$. For the diagonal approximation at $k = 1$, the relevant gap is $g_1 = \sigma_1 - \sigma_2 = 10^{-3}$ and $\mu_1 = \sqrt{2} \cdot 10^{-5}$, so




Figure 6. *The windowed Li-Li bound.* This bound states that in the windowed approximation applied to a window $\mathcal{W}$, the error is quadratic in the norm of the corresponding horizontally or vertically adjacent blocks and inversely proportional to the gap in the spectrum of the window to neighbouring windows.

the diagonal Li-Li bound evaluates as approximately $\mu_1^2/g_1 = 2 \times 10^{-7}$, matching the true error $|\sigma_1(M) - m_{11}| \approx 10^{-7}$. For the windowed approximation with $\mathcal{W} = \{1, 2\}$, the within-cluster entry m_{12} now sits inside W , the window gap is the much larger $\omega_1 = \sigma_1 - \sigma_3 = 2$, and $\mu = 10^{-5}$ is the spectral norm of either off-window block. The windowed Li-Li bound evaluates as approximately $\mu^2/\omega_1 = 5 \times 10^{-11}$, and the windowed approximation has true error $|\sigma_1(M) - \tilde{\sigma}_1| \approx 10^{-12}$, five orders of magnitude tighter than the diagonal approximation.

Proof of Theorem 3.6. The idea is to apply Li-Li to the block partition (W, C) of M , where W is the window and C is the non-window block. The majority of this proof is dedicated to showing that $\tilde{\sigma}_k$ is at position k in the merged singular values of W and C , which lets us derive this result in terms of the window gap ω_k by using the Li-Li gap η_k .

Let P be a permutation matrix reordering rows and columns so that those in $\mathcal{W}$ come first, and let P_W, P_C be the matrices selecting and permuting the window rows and columns respectively. Then

$$PMP^\top = \begin{pmatrix} W & F_1 \\ F_2 & C \end{pmatrix},$$

where $C := M[\overline{\mathcal{W}}, \overline{\mathcal{W}}]$ is the submatrix formed by the non- $\mathcal{W}$ indices, and $F_1 = P_W F P_C^\top$ and $F_2 = P_C F P_W^\top$. Since Σ is diagonal, the off-window blocks of M equal the corresponding off-window blocks of F , so

$$\|F_1\| = \|M[\mathcal{W}, \overline{\mathcal{W}}]\|, \quad \|F_2\| = \|M[\overline{\mathcal{W}}, \mathcal{W}]\|,$$

and $\mu = \max(\|F_1\|, \|F_2\|)$ in the notation of Theorem 2.8. Within each diagonal block, Weyl's inequality (Theorem 2.6) gives a loose bound on the singular values. Write $W = \Sigma_W + F_W$ where $\Sigma_W := \text{diag}(\sigma_p, \dots, \sigma_{p+\ell-1})$ and $F_W = P_W F P_W^\top$. Since $\|F_W\| \leq \delta$,

$$|\tilde{\sigma}_k - \sigma_k| \leq \delta \quad \text{for } k \in \mathcal{W}. \quad (1)$$

Likewise, write $C = \Sigma_C + F_C$, where Σ_C is the diagonal matrix collecting the entries $\sigma_{q_1}, \dots, \sigma_{q_{n-\ell}}$ enumerating the complement indices in decreasing order, so $\sigma_{q_r} = \sigma_r$ for $r \leq p-1$ and $\sigma_{q_r} = \sigma_{r+\ell}$ for $r \geq p$. Then $\|F_C\| \leq \delta$ gives

$$|\sigma_r(C) - \sigma_{q_r}| \leq \delta \quad \text{for } r = 1, \dots, n - \ell. \quad (2)$$

For any $r \in \{1, \dots, n - \ell\}$ and $k \in \mathcal{W}$, the triangle inequality with (1) and (2) gives

$$|\tilde{\sigma}_k - \sigma_r(C)| \geq |\sigma_k - \sigma_{q_r}| - 2\delta.$$

Since the singular values are decreasing, the nearest singular value to σ_k outside the window is either σ_{p-1} or $\sigma_{p+\ell}$, so $\min_r |\sigma_k - \sigma_{q_r}| = \omega_k$, giving

$$\min_{1 \leq r \leq n-\ell} |\tilde{\sigma}_k - \sigma_r(C)| \geq \omega_k - 2\delta. \quad (\star)$$

Let

$$\tilde{M} = \begin{pmatrix} W & \\ & C \end{pmatrix},$$

be the block-diagonal matrix obtained by zeroing out F_1 and F_2 . Its singular values are the merged and sorted singular values of W and C . We claim that $\sigma_k(\tilde{M}) = \tilde{\sigma}_k$ for $k \in \mathcal{W}$. For $r \leq p-1$, by (1) and (2) we have

$$\begin{aligned} \sigma_r(C) &\geq \sigma_{p-1} - \delta \\ &= \sigma_k + (\sigma_{p-1} - \sigma_k) - \delta \\ &\geq \sigma_k + \omega_k - \delta && (\text{since } p-1 \notin \mathcal{W}) \\ &> \sigma_k + \delta \\ &\geq \tilde{\sigma}_k \end{aligned}$$

and by a symmetric argument for all $r \geq p$,

$$\begin{aligned} \sigma_r(C) &\leq \sigma_{p+\ell} + \delta \\ &= \sigma_k - (\sigma_k - \sigma_{p+\ell}) + \delta \\ &\leq \sigma_k - \omega_k + \delta && (\text{since } p+\ell \notin \mathcal{W}) \\ &< \sigma_k - \delta \\ &\leq \tilde{\sigma}_k. \end{aligned}$$

So C contributes exactly $p-1$ values strictly above $\tilde{\sigma}_k$ and the remainder strictly below. Combined with W 's contribution (at most $k-p$ singular values strictly above $\tilde{\sigma}_k$, and at least $k-p+1$ values $\geq \tilde{\sigma}_k$), the merged list has at most $k-1$ values strictly above $\tilde{\sigma}_k$ and at least k values $\geq \tilde{\sigma}_k$, which implies $\sigma_k(\tilde{M}) = \tilde{\sigma}_k$.

Now we apply the Li-Li bound (Theorem 2.8) to $PMP^\top$ with the block partition (W, C) at index k ,

$$|\sigma_k(M) - \sigma_k(\tilde{M})| \leq \frac{2\mu^2}{\eta_k + \sqrt{\eta_k^2 + 4\mu^2}},$$

where, in the notation of Theorem 2.8, η_k is the minimum distance from $\sigma_k(\tilde{M})$ to the spectrum of the block that $\sigma_k(\tilde{M})$ did not originate from. To determine η_k concretely, note that, by the argument above, $\sigma_k(\tilde{M}) = \tilde{\sigma}_k$, originating from the W block, so

$$\eta_k = \min_r |\tilde{\sigma}_k - \sigma_r(C)|.$$

This is a quantity in the singular values of the submatrices W and C , whereas ω_k is defined in terms of the singular values of A itself. By (1), (2), and ($\star$), the two are related by $\eta_k \geq \omega_k - 2\delta$. Substituting this into the Li-Li bound together with the core observation $\sigma_k(M) = \sigma_k(A + E)$, we obtain

$$|\sigma_k(A + E) - \tilde{\sigma}_k| \leq \frac{2\mu^2}{(\omega_k - 2\delta) + \sqrt{(\omega_k - 2\delta)^2 + 4\mu^2}}.$$

For the unconditional bound, when $\omega_k > 2\delta$, the conditional bound and $\mu \leq \delta$ give

$$|\sigma_k(A + E) - \tilde{\sigma}_k| \leq \frac{2\delta^2}{(\omega_k - 2\delta) + \sqrt{(\omega_k - 2\delta)^2 + 4\delta^2}} \leq \frac{2\delta^2}{\omega_k},$$

since $\sqrt{(\omega_k - 2\delta)^2 + 4\delta^2} \geq 2\delta$, so the denominator is at least $(\omega_k - 2\delta) + 2\delta = \omega_k$. When $\omega_k \leq 2\delta$, (1) and Weyl applied to $A + E$ gives

$$|\sigma_k(A + E) - \tilde{\sigma}_k| \leq 2\delta \leq 2\delta \cdot \frac{2\delta}{\omega_k} \leq \frac{4\delta^2}{\omega_k}.$$

□

We note the special case of a 1×1 window, which is the diagonal Li-Li bound of Theorem 3.1.

Proof of Theorem 3.1. Apply Theorem 3.6 with $\mathcal{W} = \{k\}$, so that $p = k$ and $\ell = 1$. The window submatrix is the scalar $W = (m_{kk})$ with $\tilde{\sigma}_k = \sigma_1(W) = m_{kk}$ (using the hypothesis $m_{kk} > 0$), and the window boundary gap reduces to

$$\omega_k = \min_{r \neq k} |\sigma_k - \sigma_r| = g_k.$$

The off-window blocks reduce to the off-diagonal k th row and column of M , whose spectral norms are their Euclidean norms, so $\mu = \mu_k$. The windowed bounds therefore give

$$|\sigma_k(A + E) - m_{kk}| \leq \frac{2\mu_k^2}{(g_k - 2\delta) + \sqrt{(g_k - 2\delta)^2 + 4\mu_k^2}},$$

when $\delta < g_k/2$, and the unconditional form $|\sigma_k(A + E) - m_{kk}| \leq 4\delta^2/g_k$ when $g_k > 0$. □

3.2.3 Windowed per-pair bound via local diagonalisation

In this section, we extend the per-pair sum bound to the windowed approximation. To do this, we note that a local SVD about a window $W = U_W \Sigma_W V_W^\top$ extends to an orthogonal transformation on all of M . This gives a new matrix $M_{\text{rot}} = \tilde{U}^\top M \tilde{V}$, which has the same singular values as M but is locally diagonal within the window, and hence we may omit within-window entries inside the per-pair sum. The cost is a stronger set of assumptions on the size of the perturbation.

Theorem 3.7 (Windowed per-pair bound). *Let $M \in \mathbb{R}^{n \times n}$ have strictly decreasing positive diagonal entries $m_{11} > m_{22} > \dots > m_{nn} > 0$, and write $M = D + F_{\text{off}}$ where $D := \text{diag}(M)$ and $F_{\text{off}} := M - D$. For each $i \in \{1, \dots, n\}$, set*

$$g_i^M := \min_{j \neq i} |m_{ii} - m_{jj}|, \quad \hat{g}_i^M := \min(g_i^M, 2m_{ii}).$$

Fix a contiguous window $\mathcal{W} = \{p, p+1, \dots, p+\ell-1\} \subseteq \{1, \dots, n\}$, and let $W := M[\mathcal{W}, \mathcal{W}]$ be the corresponding submatrix. For each $k \in \mathcal{W}$, write

$$\tilde{\sigma}_k := \sigma_{k-p+1}(W),$$

for the $(k-p+1)$ -st singular value of W , viewed as the approximation to $\sigma_k(M)$ produced by the local SVD of the window.

If $\|F_{\text{off}}\| < \hat{g}_i^M/10$ for every $i \in \mathcal{W}$, then for every $k \in \mathcal{W}$,

$$|\sigma_k(M) - \tilde{\sigma}_k| \leq \frac{10}{9} \sum_{j \notin \mathcal{W}} \frac{\|M[\mathcal{W}, j]\|^2 + \|M[j, \mathcal{W}]\|^2}{|m_{kk} - m_{jj}|}.$$

Corollary 3.8 (Windowed per-pair bound for the SVD). *Let $A = U \Sigma V^\top \in \mathbb{R}^{n \times n}$ have singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$, let $E \in \mathbb{R}^{n \times n}$ be a perturbation with $\delta := \|E\|$, and write $M := \Sigma + F$ where $F := U^\top E V$. For each $i \in \{1, \dots, n\}$, set*

$$g_i := \min_{j \neq i} |\sigma_i - \sigma_j|, \quad \hat{g}_i := \min(g_i, 2\sigma_i).$$

Fix a contiguous window $\mathcal{W} = \{p, p+1, \dots, p+\ell-1\} \subseteq \{1, \dots, n\}$, let $W := M[\mathcal{W}, \mathcal{W}]$, and for each $k \in \mathcal{W}$ write $\tilde{\sigma}_k := \sigma_{k-p+1}(W)$.

If $\delta < \hat{g}_i/22$ for every $i \in \mathcal{W}$ and $m_{11} > m_{22} > \dots > m_{nn} > 0$, then for every $k \in \mathcal{W}$,

$$|\sigma_k(A + E) - \tilde{\sigma}_k| \leq \frac{10}{9} \sum_{j \notin \mathcal{W}} \frac{\|M[\mathcal{W}, j]\|^2 + \|M[j, \mathcal{W}]\|^2}{|m_{kk} - m_{jj}|}.$$

This bound has almost the same character as Corollary 3.5, but there are some subtleties in how the entries of the perturbation are weighted. Dropping the constant and rewriting the right-hand side, we obtain

$$\sum_{i \in \mathcal{W}} \sum_{j \notin \mathcal{W}} \frac{m_{ij}^2}{|m_{kk} - m_{jj}|} + \sum_{i \notin \mathcal{W}} \sum_{j \in \mathcal{W}} \frac{m_{ij}^2}{|m_{kk} - m_{ii}|}$$




Figure 7. *The windowed per-pair bound.* One interpretation of Theorem 3.7 is that we score the off-window entries in each row of $\mathcal{W}$ (blue) by $m_{ij}^2/|m_{kk} - m_{jj}|$, and the off-window entries in each column of $\mathcal{W}$ (green) by $m_{ij}^2/|m_{kk} - m_{ii}|$, i.e. the squared magnitude of each entry divided by the gap from m_{kk} to the diagonal entry sitting outside $\mathcal{W}$. The error in the windowed approximation is then the sum of these contributions from both bands.

The first sum is over the off-window entries in the rows of $\mathcal{W}$, and the second sum is over the entries in the columns. When $\mathcal{W} = \{k\}$, these sums collapse together into the diagonal per-pair bound. For larger windows, the within-window terms disappear as the local SVD has zeroed them, but every row $i \in \mathcal{W}$, not only $i = k$, now contributes to the first sum (and symmetrically for the columns in the second). Each of these contributions is still weighted by $|m_{kk} - m_{jj}|$ rather than $|m_{ii} - m_{jj}|$, so this bound is not exactly the sum of the terms in the diagonal per-pair bound.

The hypotheses of this theorem are stronger than those of the diagonal per-pair bound, which only requires $\delta < \hat{g}_k/10$ at the specific index of interest. This bound requires that $\delta < \hat{g}_i/22$ for every $i \in \mathcal{W}$. This is because the local diagonalisation replaces the diagonal of M on the window with the singular values of W , which can shift each within-window entry by up to $\|F_{\text{off}}\|$. For the new diagonal of M_{rot} to remain strictly decreasing, each within-window gap of M must be large enough so that the spectrum is not permuted, for which $\delta < \hat{g}_i/22$ for every $i \in \mathcal{W}$ is a sufficient condition. The addition of the $10/9$ factor before the sum is the penalty for writing the bound in terms of $|m_{ii} - m_{jj}|$ rather than $|m_{ii}^{\text{rot}} - m_{jj}^{\text{rot}}|$. For this reason we do not exactly recover the diagonal per-pair bound when $\ell = 1$.

To illustrate the advantage over the windowed Li-Li bound (Theorem 3.6), we return to the example of Section 3.2.2,

$$M = \begin{pmatrix} 100 & 10^{-5} & 0 & 10^{-5} \\ 10^{-5} & 100 - 10^{-3} & 0 & 0 \\ 0 & 0 & 98 & 0 \\ 10^{-5} & 0 & 0 & 1 \end{pmatrix}, \quad \mathcal{W} = \{1, 2\}.$$

The per-pair bound sees that the only non-zero off-window contribution at $k = 1$ is

at $j = 4$, paired with the much larger gap $|m_{11} - m_{44}| = 99$ rather than the window-boundary gap $\omega_1 = 2$, giving a bound of approximately $\frac{10}{9} \cdot 2(10^{-5})^2/99 \approx 2.2 \times 10^{-12}$. The true error is $|\sigma_1(M) - \tilde{\sigma}_1| \approx 10^{-12}$, so the per-pair bound is roughly 22 times tighter than Li-Li, and within a factor of ~ 2 of the truth.

Proof of Theorem 3.7. Let $W = U_W \Sigma_W V_W^\top$ be an SVD of W , and define block-diagonal orthogonal matrices $\tilde{U}, \tilde{V} \in \mathbb{R}^{n \times n}$ that act as U_W, V_W on $\mathcal{W}$ and as the identity on its complement. Set $M_{\text{rot}} := \tilde{U}^\top M \tilde{V}$. By orthogonal invariance, $\sigma_k(M_{\text{rot}}) = \sigma_k(M)$ for all k , and M_{rot} is diagonal on $\mathcal{W} \times \mathcal{W}$ with diagonal entries $\tilde{\sigma}_p, \dots, \tilde{\sigma}_{p+\ell-1}$. We verify that the conditions are sufficient to apply Theorem 3.4 to M_{rot} at index $k \in \mathcal{W}$.

Write $F_{\text{rot}} := M_{\text{rot}} - \text{diag}(M_{\text{rot}})$ and define $g_k^{\text{rot}} := \min_{j \neq k} |m_{kk}^{\text{rot}} - m_{jj}^{\text{rot}}|$ and $\hat{g}_k^{\text{rot}} := \min(g_k^{\text{rot}}, 2m_{kk}^{\text{rot}})$. We have $\|F_{\text{rot}}\| \leq 2\|F_{\text{off}}\|$, since F_{rot} differs from $\tilde{U}^\top F_{\text{off}} \tilde{V}$ only in the window block. For the gap, Weyl on W and the reverse triangle inequality give $g_k^{\text{rot}} \geq g_k^M - 2\|F_{\text{off}}\|$ and $m_{kk}^{\text{rot}} \geq m_{kk} - \|F_{\text{off}}\|$. Hence $\hat{g}_k^{\text{rot}} \geq \hat{g}_k^M - 2\|F_{\text{off}}\|$, and then

$$\|F_{\text{rot}}\| \leq 2\|F_{\text{off}}\| < \frac{2}{10}\hat{g}_k^M < \frac{1}{4}(\hat{g}_k^M - 2\|F_{\text{off}}\|) \leq \frac{1}{4}\hat{g}_k^{\text{rot}}.$$

To see that M_{rot} has a strictly decreasing positive diagonal, first note that by Weyl's inequality applied to W ,

$$|\sigma_j(W) - m_{p+j-1, p+j-1}| \leq \|F_{\text{off}}\| \quad \text{for each } j = 1, \dots, \ell.$$

Then within $\mathcal{W}$, consecutive entries of M_{rot} 's diagonal are separated. Writing $\Delta_j := m_{p+j-1, p+j-1} - m_{p+j, p+j}$ for the consecutive gaps in M 's diagonal (which are positive by the hypothesis), Weyl applied twice to W gives

$$\sigma_j(W) - \sigma_{j+1}(W) \geq \Delta_j - 2\|F_{\text{off}}\|,$$

for each $j = 1, \dots, \ell - 1$. Since $\|F_{\text{off}}\| < g_{p+j-1}^M/10 \leq \Delta_j/10$ by hypothesis, this is at least $\frac{4}{5}\Delta_j > 0$.

At the window boundaries, Weyl gives $\sigma_1(W) \leq m_{pp} + \|F_{\text{off}}\|$ and $\sigma_\ell(W) \geq m_{p+\ell-1, p+\ell-1} - \|F_{\text{off}}\|$. Since $g_p^M \leq m_{p-1, p-1} - m_{pp}$ when $p > 1$, the hypothesis at $i = p$ gives $\|F_{\text{off}}\| < g_p^M/10 \leq (m_{p-1, p-1} - m_{pp})/10$. Hence $\sigma_1(W) < m_{p-1, p-1}$. Symmetrically, $g_{p+\ell-1}^M \leq m_{p+\ell-1, p+\ell-1} - m_{p+\ell, p+\ell}$ when $p + \ell - 1 < n$, and the hypothesis at $i = p + \ell - 1$ gives $\sigma_\ell(W) > m_{p+\ell, p+\ell}$. Hence the within-window subset of M_{rot} 's diagonal sits strictly between $m_{p-1, p-1}$ and $m_{p+\ell, p+\ell}$. Since the singular values of W are strictly ordered and $\sigma_\ell(W) > m_{p+\ell, p+\ell}$, we also have the required positivity.

Applying Theorem 3.4 to M_{rot} at index k gives

$$|\sigma_k(M) - \tilde{\sigma}_k| \leq \sum_{j \neq k} \frac{(m_{kj}^{\text{rot}})^2 + (m_{jk}^{\text{rot}})^2}{|m_{kk}^{\text{rot}} - m_{jj}^{\text{rot}}|},$$

using $\sigma_k(M_{\text{rot}}) = \sigma_k(M)$ and $m_{kk}^{\text{rot}} = \tilde{\sigma}_k$. For $j \in \mathcal{W} \setminus \{k\}$, the terms are zero as M_{rot} is diagonal on $\mathcal{W} \times \mathcal{W}$. For $j \notin \mathcal{W}$, $m_{jj}^{\text{rot}} = m_{jj}$ and the orthonormality of U_W and V_W gives

$$(m_{kj}^{\text{rot}})^2 \leq \|M[\mathcal{W}, j]\|^2 \quad \text{and} \quad (m_{jk}^{\text{rot}})^2 \leq \|M[j, \mathcal{W}]\|^2.$$

For the denominator, Weyl on W gives $|\tilde{\sigma}_k - m_{kk}| \leq \|F_{\text{off}}\|$, and the reverse triangle inequality gives

$$\begin{aligned} |\tilde{\sigma}_k - m_{jj}| &\geq |m_{kk} - m_{jj}| - \|F_{\text{off}}\| \\ &\geq \frac{9}{10} |m_{kk} - m_{jj}|, \end{aligned}$$

where the second line uses $\|F_{\text{off}}\| < g_k^M/10 \leq |m_{kk} - m_{jj}|/10$. Hence $1/|\tilde{\sigma}_k - m_{jj}| \leq (10/9)/|m_{kk} - m_{jj}|$, and so

$$|\sigma_k(M) - \tilde{\sigma}_k| \leq \frac{10}{9} \sum_{j \notin \mathcal{W}} \frac{\|M[\mathcal{W}, j]\|^2 + \|M[j, \mathcal{W}]\|^2}{|m_{kk} - m_{jj}|},$$

as required. $\square$

Proof of Corollary 3.8. Write $F_{\text{off}} := F - \text{diag}(F) = M - \text{diag}(M)$. The diagonal entry of any matrix is bounded in modulus by its spectral norm, so $|f_{ii}| \leq \|F\| = \delta$ for each i , and hence $\|F_{\text{off}}\| \leq \|F\| + \|\text{diag}(F)\| \leq 2\delta$. For each $i \in \mathcal{W}$, two applications of the diagonal bound via the reverse triangle inequality give

$$|m_{ii} - m_{jj}| \geq |\sigma_i - \sigma_j| - 2\delta \quad \text{for each } j \neq i,$$

so $g_i^M \geq g_i - 2\delta$. A single application gives $m_{ii} \geq \sigma_i - \delta$. By definition of $\hat{g}_i$, the hypothesis $\delta < \hat{g}_i/22$ is equivalent to $\delta < g_i/22$ and $\delta < \sigma_i/11$. The first gives $g_i^M > 20\delta \geq 10\|F_{\text{off}}\|$. The second gives $m_{ii} > 10\delta$, so $2m_{ii} > 20\delta \geq 10\|F_{\text{off}}\|$. Hence $\hat{g}_i^M := \min(g_i^M, 2m_{ii}) > 10\|F_{\text{off}}\|$, i.e. $\|F_{\text{off}}\| < \hat{g}_i^M/10$. Now Theorem 3.7 applied to M , together with the identity $\sigma_k(A + E) = \sigma_k(M)$ of Corollary 2.5, gives the claimed bound. $\square$

3.3 The banded approximation




Figure 8. *The banded approximation.* We approximate $\sigma_k(M)$ by the SVD of M truncated to bandwidth b .

We conclude this section with a brief discussion of the banded approximation. The *banded approximation* retains the entries of M within a band of width b around the diagonal:

$$M_{ij}^{(b)} := \begin{cases} m_{ij} & |i - j| \leq b, \\ 0 & |i - j| > b \end{cases} \quad \tilde{\sigma}_k(M) := \sigma_k(M^{(b)}).$$

We do not have theoretical bounds for this approximation and the literature on off-band perturbations like this is sparser (see for example Nakatsukasa [17], which studies a related setting of eigenvalue perturbation bounds for block tridiagonal matrices undergoing a block tridiagonal perturbation). We conjecture a result like Corollary 3.5 where the sum is over pairs with $|k - j| > b$, but we cannot use the same technique as in Corollary 3.8 as the local diagonalisation in this case affects all off-band entries of M .

4 Experiments

Table 2. *Spectrum notation.* In several experiments, we see how the error changes depending on the spectrum of the matrix under consideration. We consider linear, geometric, and clustered spectra. Where useful we include plots of the singular values.

Notation	Definition
linear	$\sigma_k = n - k + 1$
geometric $_{r,A}$	$\sigma_k = 1 + A \cdot r^{k-1}$ with ratio $r \in (0, 1)$ and amplitude A , we write geometric $_r$ when $A = 1$.
clustered $_{s,g}^c$	c equally sized clusters with spread within each cluster s and a gap of g between consecutive clusters. We write clustered $_{s,g}$ when $c = 2$.

Each plot below is the mean of 20 trials. We arrange spectra left-to-right and our approximations top-to-bottom.

4.1 The approximations in practice

We characterise the empirical behaviour of each approximation as we vary the perturbation size, the spectrum, and the dimension.




Figure 9. *Error vs δ .* $n = 200$, $\delta \in [10^{-4}, 10^{-1}]$, and linear, $\text{geometric}_{r=0.99}$, $\text{clustered}_{s=0.1, g=1}$ spectra. We plot the max error over all singular values for the diagonal, windowed (window size $\ell \times \ell$), and banded (bandwidth b) approximations. All methods follow $O(\delta^2)$, and increasing the window size or bandwidth gives a modest improvement, except for the largest window size in the clustered spectrum case which is large enough to capture the clustered spectrum structure.




Figure 10. *Per-index error.* $n = 50$, $\delta = 10^{-2}$, and linear, $\text{geometric}_{r=0.99}$, $\text{clustered}_{s=0.1, g=0.9}$ spectra. We plot the per-index error of each approximation. Note the oscillating pattern in the windowed approximation, which corresponds to the middle index of each window having smallest error.




Figure 11. *Error vs gap.* $n = 100$, $\delta = 10^{-2}$, spectrum $[99 + g_1, 99, 98, \dots, 1]$ with $g_1 \in [10^{-4}, 10^2]$. We plot the error in the first singular value for each approximation.




Figure 12. *Error vs n .* $\delta = 10^{-2}$, $n \in [30, 1000]$, and linear, $\text{geometric}_{r=0.99}$, $\text{clustered}_{s=0.1, g=0.9}$ spectra. For the linear spectrum the error decreases as $O(\delta^2 \ln n/n)$, matching the scaling derived for this spectrum from Corollary 3.5. For the geometric spectrum, the gap decays fast enough with n that the denominator contributes a factor of approximately r^n . For the clustered spectrum, the error stays approximately constant since the gap shrinks with n .




Figure 13. *Time vs n .* $\delta = 10^{-3}$, $n \in [100, 5000]$, and a linear spectrum. We plot the best of 5 wall-clock measurements (including the cost of forming M) to compute the approximate singular values by the diagonal approximation, windowed approximation and full SVD. We omit the timings for the banded approximation due to difficulties implementing an efficient $O(n^2b)$ banded SVD algorithm.




Figure 14. *Accuracy vs time.* $n = 500$, $\delta = 10^{-3}$, and linear, $\text{geometric}_{r=0.9}$, $\text{clustered}_{s=0.1,g=0.9}$ spectra over 5 trials. We plot $(\text{time}, \max_k \text{error})$ for the diagonal and windowed approximations while varying the window size $\ell \in \{2, 5, 10, 15, 25, 40, 60, 90, 140, 200\}$. The full SVD sits in the bottom-right corner. Increasing the window size is slower and only gives modest improvements in accuracy.

4.2 Tightness of the bounds

In the following plots, we verify that the hypotheses of each of our bounds hold before applying them, with the exception of Figure 19, where we see what happens to the error outside the regime where the hypothesis holds.

4.2.1 Diagonal method

We compare the diagonal Li-Li, worst-pair, and per-pair bounds against the empirical error.




Figure 15. *Per-index error and diagonal bounds.* $\delta = 10^{-3}$, and linear, $\text{geometric}_{r=0.95}$, $\text{clustered}_{s=0.2,g=1}$ spectra. We plot the per-index error for the diagonal method against the worst-pair bound (Corollary 3.3), the Li-Li bound (Theorem 3.1), and the per-pair bound (Corollary 3.5). We further plot each spectrum and the corresponding matrix $(m_{ij}^2 + m_{ji}^2)/|m_{ii} - m_{jj}|$ whose column and row sums give the per-pair bound. The worst-pair bound is very pessimistic and uniform across k . For the clustered spectrum, the per-pair bound predicts that the error should peak in the middle of each cluster, but in practice it is actually the lowest there. We attribute this to applying the triangle inequality to a particular signed sum in the proof of Theorem 3.4.




Figure 16. *Diagonal bounds compared to the error vs n .* $\delta = 10^{-3}$, $n \in [10, 200]$, and linear, $\text{geometric}_{r=0.99}$, $\text{clustered}_{s=0.5, g=2}$ spectra. We plot the max error over all singular values when using the diagonal method and compare against the worst-pair bound (Corollary 3.3), the Li-Li bound (Theorem 3.1), and the per-pair bound (Corollary 3.5). The per-pair bound is the tightest and the only bound that correctly predicts that the error should shrink with n in the linear case (see Figure 12).




Figure 17. *Per-index error and diagonal bounds for a three-cluster spectrum.* $n = 50$, $\delta = 10^{-4}$, and a three-cluster spectrum with 20, 15, 15 singular values evenly spaced over the intervals $[9, 10]$, $[4, 5]$, $[1, 2]$ respectively. We plot the per-index error and bounds as in Figure 15. Compared to Figure 15, the Li-Li bound noticeably decreases as we move to the second cluster. This is because the second two clusters are slightly smaller and so the intra-cluster gap is larger. The same per-pair concavity observed in Figure 15 appears here too.




Figure 18. *Adversarial construction against Li-Li.* An example where Li-Li can be made arbitrarily large while our bounds remain constant. Take $A = \text{diag}(1 + \epsilon, 1, 0.5)$ and $E = 10^{-6}(e_1 e_3^\top + e_3 e_1^\top)$, sweeping ϵ in the example of Theorem 3.2. As A is diagonal, $F = U^\top E V = E$, and at index 1 the off-row/column norm is $\mu_1 = 10^{-6} = \delta$. The Li-Li bound at index 1 is $4\mu_1^2/|\sigma_1 - \sigma_2| = 4\delta^2/\epsilon$, which uses the wrong gap. However, since E is only non-zero at entries $E_{1,3} = E_{3,1} = 10^{-6}$, the $|\sigma_1 - \sigma_2|$ term is dropped from both per-pair and worst-pair bounds, giving error $2\delta^2/|\sigma_1 - \sigma_3| = 2\delta^2/(\epsilon + 0.5) \approx 4\delta^2$ when ϵ is small. Hence the per-pair and worst-pair bounds are independent of ϵ but the Li-Li bound can be made arbitrarily large by sending $\epsilon \rightarrow 0$. For small enough ϵ , the worst-pair bound can be smaller than the Li-Li bound despite the n^2 prefactor.




Figure 19. *Behaviour of the bounds outside hypotheses.* $n = 30$ and a linear spectrum $10^{-3} \cdot \text{diag}(30, 29, \dots, 1)$. We plot the max error across the whole spectrum against the per-pair bound (Corollary 3.5), the worst-pair bound (Corollary 3.3), the unconditional Li-Li bound (Theorem 3.1), and the Weyl bound (Theorem 2.6). This spectrum has $g_i = \hat{g}_i = 10^{-3}$ for all i . We shade three regions. In the first, the per-pair hypothesis ($\delta < \hat{g}_i/10$) holds. In the second, the tight Li-Li hypothesis ($\delta < g_i/2$) holds. In the final region neither holds. The worst-pair bound has no hypotheses. After $\delta > g_i/2$ the Weyl bound dominates, and the error trends as $O(\delta)$.

4.2.2 Windowed method

We compare the windowed Li-Li and windowed per-pair bounds across window sizes and spectra.




Figure 20. Bounds for the windowed method vs n . $\delta = 10^{-3}$, $n \in [30, 1000]$ varies, and linear, geometric $_{r=0.999, A=100}$, clustered $_{s=15, g=20}$ spectra. Rows are window sizes ℓ . Within each row we plot the max error against the windowed Li-Li bound (Theorem 3.6) and windowed per-pair bound (Corollary 3.8), with the results for the other rows shown faintly. Accuracy increases slightly with window size, but as ℓ grows the per-pair bound actually loosens, since the sum of off-window entries grows faster than the benefit from zeroing the within-window sum. The Li-Li bound stays consistent across all window sizes.




Figure 21. *Per-index error and bounds for the windowed method.* $n = 30$, $\delta = 10^{-4}$, spectra linear, geometric $_{r=0.95}$, clustered $_{s=0.2, g=1.0}$. Rows are window sizes ℓ . Within each cell we plot the per-index error against the windowed Li-Li bound (Theorem 3.6) and windowed per-pair bound (Corollary 3.8). Alternating shaded bands mark the placement of windows. The bound is smallest in the middle of the window, where the boundary gap ω_j is largest, and increases towards the window edges, where ω_j shrinks to the gap with the nearest singular value outside the window.




Figure 22. *Windowed method on a three-cluster spectrum.* $n = 50$, $\delta = 10^{-5}$, three-cluster spectrum with 20, 20, 10 singular values evenly spaced over the intervals $[9, 10]$, $[4, 5]$, $[1, 2]$ respectively. Within each cell we plot the per-index error against the windowed Li-Li bound (Theorem 3.6) and windowed per-pair bound (Corollary 3.8). Li-Li can be tighter than the per-pair bound for windows because the per-pair numerator $\|M[\mathcal{W}, j]\|^2 + \|M[j, \mathcal{W}]\|^2$ summed over j is the squared Frobenius norm of the blocks, which can exceed the squared spectral norm μ^2 used in the Li-Li bound.

5 Conclusion

We have proved and empirically studied a hierarchy of $O(\mu^2/\text{gap})$ error bounds for the diagonal and windowed approximations to the singular values of $A + E$. Each bound in the hierarchy uses finer information about how the perturbation and spectrum of $A + E$ enter the estimate, through the minimum gap (Theorem 3.1), the worst pairwise gap (Theorem 3.2), and finally the per-pair gap (Theorem 3.4). Our numerical experiments have demonstrated the effectiveness of our bounds experimentally.

There are several avenues that could be explored in a future work.

Rectangular and complex matrices. Our results as stated apply to real square matrices. While similar statements for rectangular and complex matrices seem straightforward, there are some subtleties that would need to be addressed. For example, the Jordan-Wielandt construction for a rectangular matrix introduces additional 0 eigenvalues, which would need to be treated carefully in our proof of Theorem 3.4.

Hermitian eigenvalue problems. We have focused on perturbation bounds for singular values. Many perturbation results for singular values are corollaries of more general theorems about Hermitian eigenvalue problems. It may be the case that there is an eigenvalue analogue of Theorem 3.4 with a simpler proof that avoids the need to carry the Jordan-Wielandt matrix through the analysis, and from which we may derive our singular value bound as a consequence.

Extensions of the per-pair bound. We see two possible approaches to refine the analysis. First, we noted in the proof of Theorem 3.4 a step which leads to a nice symmetric form for the final bound but introduces unnecessary looseness. We conjecture that a version of the per-pair bound of the form

$$|\sigma_k(M) - m_{kk}| \lesssim \left| \sum_{j \neq k} \frac{(m_{jk} + m_{kj})^2}{2|m_{kk} - m_{jj}|} + \sum_{j \neq k} \frac{(m_{jk} - m_{kj})^2}{2(m_{kk} + m_{jj})} \right|$$

may hold. We also note that it may be possible to work directly with the window block matrix in the Schur complement step, from which a tighter windowed per-pair bound may follow.

Singular vectors. We left a precise treatment of the singular vectors out of the scope of this dissertation, but approximate singular vectors arise naturally in the windowed and banded approximations. Further work could study these vectors in more detail and derive Davis-Kahan-style bounds of the form $\sin \theta \leq O(\mu/\text{gap})$ where θ is the angle between the approximate singular vectors, although this would require different techniques than those we have used here.

References

- [1] The Schur complement and its applications, 2005. URL <https://doi.org/10.1007/b105056>.
- [2] Noah Amsel, David Persson, Christopher Musco, and Robert M. Gower. The Polar Express: Optimal Matrix Sign Methods and Their Application to the Muon Algorithm, May 2025. URL <https://arxiv.org/abs/2505.16932>.
- [3] Matthew Brand. Fast low-rank modifications of the thin singular value decomposition. *Linear Algebra and its Applications*, 415(1):20–30, 2006. ISSN 00243795. doi: 10.1016/j.laa.2005.07.021. URL <https://doi.org/10.1016/j.laa.2005.07.021>.
- [4] James R. Bunch and Christopher P. Nielsen. Updating the singular value decomposition. *Numerische Mathematik*, 31(2):111–129, 1978. ISSN 0945-3245. doi: 10.1007/BF01397471. URL <https://doi.org/10.1007/BF01397471>.
- [5] James R. Bunch, Christopher P. Nielsen, and Danny C. Sorensen. Rank-one modification of the symmetric eigenproblem. *Numerische Mathematik*, 31(1):31–48, 1978. ISSN 0945-3245. doi: 10.1007/BF01396012. URL <https://doi.org/10.1007/BF01396012>.
- [6] Chandler Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970. ISSN 00361429. doi: 10.1137/0707001. URL <https://doi.org/10.1137/0707001>.
- [7] Gene H. Golub. Some modified matrix eigenvalue problems. *SIAM Review*, 15(2):318–334, 1973. ISSN 00361445, 10957200. doi: 10.1137/1015032. URL <https://doi.org/10.1137/1015032>.
- [8] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 4th edition, 2013.
- [9] Ming Gu and Stanley C. Eisenstat. A stable and efficient algorithm for the rank-one modification of the symmetric eigenproblem. *SIAM Journal on Matrix Analysis and Applications*, 15(4):1266–1276, 1994. ISSN 0895-4798, 1095-7162. doi: 10.1137/S089547989223924X. URL <https://doi.org/10.1137/S089547989223924X>.
- [10] C. G. J. Jacobi. Über ein leichtes verfahren die in der theorie der säcularstörungen vorkommenden gleichungen numerisch aufzulösen. *Journal für die reine und angewandte Mathematik*, 30:51–94, 1846. doi: 10.1515/crll.1846.30.51. URL <https://doi.org/10.1515/crll.1846.30.51>.
- [11] Camille Jordan. Mémoire sur les formes bilinéaires. *Journal de Mathématiques Pures et Appliquées*, 19:35–54, 1874. URL https://www.numdam.org/item/JMPA_1874_2_19__35_0/.

- [12] Maseim Kenmoe, Ronald Kriemann, Matteo Smerlak, and Anton S. Zadorin. A fast iterative algorithm for near-diagonal eigenvalue problems. *SIAM Journal on Matrix Analysis and Applications*, 43(4):1573–1598, 2022. doi: 10.1137/21M1401474. URL <https://doi.org/10.1137/21M1401474>.
- [13] E. G. Kogbetliantz. Solution of linear equations by diagonalization of coefficients matrix. *Quarterly of Applied Mathematics*, 13(2):123–132, 1955. ISSN 0033569X, 15524485. doi: 10.1090/qam/88795. URL <https://doi.org/10.1090/qam/88795>.
- [14] Chi-Kwong Li and Ren-Cang Li. A note on eigenvalues of perturbed hermitian matrices. *Linear Algebra and its Applications*, 395:183–190, 2005. ISSN 0024-3795. doi: 10.1016/j.laa.2004.08.026. URL <https://doi.org/10.1016/j.laa.2004.08.026>.
- [15] Roy Mathias. Quadratic residual bounds for the hermitian eigenvalue problem. *SIAM Journal on Matrix Analysis and Applications*, 19(2):541–550, 1998. ISSN 0895-4798, 1095-7162. doi: 10.1137/S0895479896310536. URL <https://doi.org/10.1137/S0895479896310536>.
- [16] L. Mirsky. Symmetric gauge functions and unitarily invariant norms. *The Quarterly Journal of Mathematics*, 11(1):50–59, 1960. ISSN 0033-5606. doi: 10.1093/qmath/11.1.50. URL <https://doi.org/10.1093/qmath/11.1.50>.
- [17] Yuji Nakatsukasa. Eigenvalue perturbation bounds for hermitian block tridiagonal matrices. *Applied Numerical Mathematics*, 62(1):67–78, 2012. ISSN 0168-9274. doi: 10.1016/j.apnum.2011.09.010. URL <https://doi.org/10.1016/j.apnum.2011.09.010>.
- [18] Yuji Nakatsukasa. Off-diagonal perturbation, first-order approximation and quadratic residual bounds for matrix eigenvalue problems. In Tetsuya Sakurai, Shao-Liang Zhang, Toshiyuki Imamura, Yusaku Yamamoto, Yoshinobu Kura-mashi, and Takeo Hoshi, editors, *Eigenvalue Problems: Algorithms, Software and Applications in Petascale Computing*, pages 233–249. Springer International Publishing, 2017. ISBN 978-3-319-62426-6. doi: 10.1007/978-3-319-62426-6_15. URL https://doi.org/10.1007/978-3-319-62426-6_15.
- [19] Nian Shao and Yuji Nakatsukasa. Accelerating preconditioned Jacobi methods via perturbation-inspired pivoting. Manuscript in preparation, 2025.
- [20] Gilbert W. Stewart and Ji-guang Sun. *Matrix Perturbation Theory*. Computer Science and Scientific Computing. Academic Press, Boston, 1990. ISBN 0126702306 9780126702309.
- [21] Ji-guang Sun. Eigenvalues of rayleigh quotient matrices. *Numerische Mathematik*, 59(1):603–614, 1991. ISSN 0945-3245. doi: 10.1007/BF01385798. URL <https://doi.org/10.1007/BF01385798>.
- [22] Per-Åke Wedin. Perturbation bounds in connection with singular value decom-

- position. *BIT Numerical Mathematics*, 12(1):99–111, 1972. ISSN 1572-9125. doi: 10.1007/BF01932678. URL <https://doi.org/10.1007/BF01932678>.
- [23] Hermann Weyl. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, 1912. ISSN 1432-1807. doi: 10.1007/BF01456804. URL <https://doi.org/10.1007/BF01456804>.