

Eigenvalue Enclosures for Definite Hermitian Matrix Pencils under Structured Low-Rank Perturbations

Authors Names	ABSTRACT
Hussein Mhmood Ali Algburi Publication date: 26 /8 /2026 Keywords: Definite Hermitian matrix pencil; eigenvalue enclosure; structured low-rank perturbation; generalized eigenvalue problem; inertia; interlacing.	This study develops an eigenvalue-enclosure framework for definite Hermitian matrix pencils subject to simultaneous structured low-rank perturbations of both coefficient matrices. After reducing a definite pair to a positive-definite denominator, the perturbations are compressed to a common low-dimensional support. A shifted global enclosure is first obtained from Loewner inequalities. A projection-aware two-sided bound is then derived from the generalized Courant–Fischer principle; its index-dependent factors measure the overlap between the update support and lower or upper invariant subspaces. A rank-interlacing guard is intersected with these bounds. For selected eigenvalues, a determinant and inertia identity expresses the updated eigenvalue count through a correction matrix whose order equals the support dimension, enabling safeguarded bisection without a complete updated eigendecomposition. The construction is extended to general definite Hermitian pairs through a real rotation and a monotone Möbius map. In deterministic experiments on 216 real and complex Hermitian pencils (20,736 ordered eigenvalues), no coverage failure was observed and the combined intervals reduced the shifted global widths by 49.6% on average. A six-eigenvalue cluster was refined from mean initial width 2.23×10^{-1} to 5.19×10^{-11}, and a localized mass–spring example produced a 73.1% mean reduction for selected modes. The method is most suitable for repeated low-rank updates when baseline spectral information is reusable and only selected eigenvalues are required.

1. Introduction

Definite generalized Hermitian eigenvalue problems are problems which involve pairs of Hermitian matrices. Some typical examples are vibration analysis, finite-element discretizations, stability assessment, constrained optimization and physical systems that depend on a parameter. These applications can have a baseline pencil solution that can be repeatedly tweaked by a small number of localized changes. These kinds of changes can manifest themselves in changes to either the stiffnesslike matrix or to the mass/metric-like matrix, and are in fact naturally represented by structured low-rank perturbations.

Let
$$L(\lambda) = A - \lambda B, \quad A = A^*, \quad B = B^*, \quad (1)$$

and let the perturbed pencil be

$$\tilde{L}(\lambda) = A + E - \lambda(B + F), \quad E = E^*, \quad F = F^*. \quad (2)$$

The key issue is more pragmatic than the above: does the eigenspectrum change continuously, and can each individual ordered perturbed eigenvalue be bound in a computable interval that is much smaller than the global normwise perturbation bound? This difference is significant if the perturbation has low rank. An overall magnitude of this movement is recorded in a global norm, while the actual movement of a selected eigenvalue also depends on the orientation of the perturbation range with respect to the pertinent invariant subspaces.

More recent literature on the subject includes a series of complementary developments, namely, definite generalized-Hermitian perturbation bounds [1]; low-rank eigenstructure and placement [2]; symmetry-preserving low-rank changes [3]; spectral analysis of pencils with Hermitian parts [4]; definite-pair eigenspace bounds [5]; verified Hermitian inclusions [6]; bounded-rank pencil theory [12]–[14]; special structured linear perturbations [10];

parametric rank-one motion [11]; structure-preserving generalized eigensolvers [12]; minimal-rank perturbations [13]; and inverse-free pencil diagonalization [14] as well as related fixed-rank, rank-one, bounded-rank, distinct-eigenvalue and parametric developments [15]–[19]. These developments lay the groundwork but fail to offer a framework that offers an update-based approach while at the same time introducing an ordered enclosure for the two-sided eigenvalues as well as also taking advantage of the low rank orientation, and lastly, narrowing down selected intervals by fewer eigenvalues to be encountered.

Research gap. While much of the literature on pencils of low rank is algebraic and aims at characterizing admissible eigenstructures, the existing definite-pair perturbation bounds are mostly global, normwise, and/or particular to certain partitions of the pencil. Verified full-spectrum methods on the other hand, do not tailor themselves specifically for repeated structured low rank changes. The area of overlap of these three topics: definiteness, low-rank geometry, and index-wise interval certification remains underdeveloped.

Research problem. When the order of Hermitian matrix and a number of Hermitian perturbations are much larger than the number of Hermitian perturbations, given a definite Hermitian pencil, compute intervals where each or a subset of the ordered perturbed eigenvalues lie and keep information on definiteness without recomputing the full perturbed generalized eigensolve for every iteration.

This paper tackles the problem in the following five contributions.

1. A shifted global enclosure is derived for simultaneous perturbations of both matrices of a positive-definite Hermitian pair.
2. A projection-aware enclosure is established. Its width depends on the overlap between the low-rank update space and lower/upper invariant subspaces associated with each ordered eigenvalue.
3. A rank-interlacing guard is proved and intersected with the two analytic enclosures.
4. A determinant-and-inertia reduction converts the perturbed eigenvalue count at a test point into the inertia of a matrix whose dimension equals the low-rank support dimension. This yields a safeguarded bisection procedure for selected eigenvalues.
5. The positive-definite formulation is extended to general definite Hermitian pairs by a real rotation and a monotone Möbius map.

This numerical study assesses these claims vis-à-vis their target objectives. The results indicated that the combined intervals had no observed coverage failure across all 216 test pencils consisting of 20,736 ordered eigenvalues, and were able to improve the global widths by 49.6% on average. A three-means reduced-inertia refinement for an eight value-cluster involved increasing the smallest of the initial eigenvalues from a typical minimum value of 2.23×10^{-1} to 5.19×10^{-11} . A localized mass spring update resulted in an average of 73.1% reduction in width for six modes reported. These are observations with finite accuracy in support of the theory and not proof of what is shown below.

The rest of the paper will be organized as follows. A review of literature in recent years has been provided in Section 2 stating the exact relationship between each stream and the present study. The objectives are presented in Section 3, along with the mathematical context. The enclosure and inertia theorems are presented in section 4. The complexity and the computational procedure is given in Section 5. The numerical experiments are reported in section 6. The results are discussed with respect to the objectives and related to the previous literature in Section 7 and the paper is summed up in Section 8.

2. Related Literature and Position of the Present Study

2.1 Definite Hermitian pair perturbation

Xiao and Li [1] established absolute perturbation bounds for partitioned generalized Hermitian positive-definite eigenvalue problems based on the blocks, and demonstrated that block information can improve absolute perturbation bounds that do not take into account the partition. In [5] Kuzmanović Ivičić, Miodragović, and Ugriča presented a $\tan\Theta$ matrix for definite matrix pairs and showed that structure can play a major role in improving estimates of their eigenspaces. All of these studies point to the main theme: geometry of a perturbation is of importance. They aim for, however, different goals than the current goal: an all-index low-rank interval framework, together with a reduced inertia count when updating both the numerator and denominator, are not encouraged.

Rump and Lange [6] offered verified error bounds for all pairs of eigenvalues of Hermitian matrices, both for close

clusters as well as for multiple eigenvalues. That study lays the groundwork for recognizing the worth of – and challenges to – rigorous inclusion strategies. While its setting is complementary to the update-oriented setting considered here, where a baseline eigendecomposition is assumed to be available and the new information lies in low-dimensional subspace, its full-spectrum verification setting is interesting.

2.2 Algebraic and structured low-rank pencil perturbations

These were Weierstrass structures and placement of eigenvalues in the presence of perturbations that are bounded and have low rank, treated by Baragaña and Roca [2] and subsequently fixed rank perturbations treated [15]. Dodig and Stošić solved rank-one perturbations of arbitrary pencils [16], carried out investigations in the bounded-rank and quasi-regular cases [7, 9, 13] and posed the perturbation problem as a completion problem. Arrows and Moon [1] proposed a similar problem to treat the properties when the number of distinct eigenvalues is bounded due to a perturbation. Answers are provided to important questions concerning existence and canonical-structure, but no numerical intervals $[l_i, u_i]$ are given in the ordered real eigenvalues of a particular definite Hermitian pair.

This work was inspired by De Terán et al. [3] which perform studies of structure-preserving low-rank perturbations, for several pencil symmetries. Eigenstructures of Hermitian pencils are discussed under the name of generic eigenstructures by de Terán, Dmytryshyn, and Dopico [8] by emphasizing on the role of rank and sign characteristics. Mehl, Mehrmann, and Wojtylak [4] studied pencils with positive semidefinite Hermitian parts and obtained information on their spectral locations. Inspired by the works mentioned above we consider an explicit preservation of Hermitian structure in the present model. The current paper asks a quantitative, ordered, and computational question after it is specified what the update is.

The behaviour of parametric rank-one perturbations (eigenvalue motion, multiplicity change, asymptotic behaviour, an interlacing result) were analyzed by Gernandt and Trunk [11]. In particular, their rank-one trajectory analysis is very much in the spirit of the rank guard employed here. There are three changes in the present framework: (i) an r general dimension of support s , (ii) possible simultaneous changes of both matrices, (iii) a diminished number of change points for finite intervals instead of a single scalar perturbation trajectory.

For some specific structured pencils, Ahmadnasab and Psarrakos [10] derived eigenvalue characterizations and location bounds in the presence of linear perturbations, including specific families of pencils motivated by control. They have shown that their special structure can result in more precise location information. The present study instead considers the definite Hermitian pair of arbitrary shape, reduced to positive-definite denominator and considers as actual support orientation.

2.3 Structure-preserving and modern pencil algorithms

Based on Hermitian perturbations, one of the authors discussed structure preserving methods for solving singular generalized eigenvalue problems and focused particularly on the preservation of having sign characteristics [12]. A randomized inverse-free diagonalization framework for general matrix pencils was developed by Demmel, Dumitriu and Schneider [14]. Such works illustrate that structure and algorithm design issues are still very much current themes at the forefront of generalized eigenvalue computation. The full idea behind the proposed enclosure method would not be to use it in place of a complete eigensolver, but it is meant as a localization update after baseline eigensolver computation only.

2.4 Comparative analysis of representative studies

The relationship of the representative recent studies to the present contribution is summarized in Table 1. The final column notes a contextual deficiency in regard to the subject under study in the current research from which the concern was generated as opposed to a deficiency within the object of the study as identified by the author.

Study	Principal contribution	Relation to the present study	Remaining need addressed here
Xiao and Li [1], 2018	Sharper absolute bounds for partitioned definite generalized Hermitian problems	Confirms that substructure can improve global estimates	General factorized low-rank updates to both matrices and ordered intervals
Baragaña and Roca [2], 2019	Weierstrass structure and eigenvalue placement under low-rank perturbations	Establishes algebraic possibilities after low-rank changes	Computable numerical enclosure for a specified updated pencil
Dodig and Stošić [16], 2020	Complete characterization of rank-one pencil perturbations	Provides canonical rank-one foundation	General rank- S quantitative intervals and reduced counting

Study	Principal contribution	Relation to the present study	Remaining need addressed here
De Terán et al. [3], 2022	Generic structure-preserving low-rank changes for symmetric/Hermitian and related pencils	Motivates Hermitian update factors	Index-wise enclosure widths for a definite pair
Mehl et al. [4], 2022	Spectral properties of pencils with positive semidefinite Hermitian parts	Connects Hermitian parts to spectral location	Update-specific two-sided ordered bounds
Kuzmanović Ivičić et al. [5], 2023	Eigenspace angle bounds for definite pairs	Shows gains from structured geometry	Eigenvalue intervals and low-rank inertia refinement
Rump and Lange [6], 2023	Verified bounds for all Hermitian eigenpairs, including clusters	Provides a rigorous verification benchmark	Repeated low-rank update specialization and target-only cost
De Terán et al. [8], 2024	Generic eigenstructures of Hermitian pencils	Clarifies rank and sign-characteristic effects	Quantitative ordered localization for a given definite pencil
Dodig and Stošić [9], 2024	Bounded-rank perturbations of regular pencils	Advances low-rank structural theory	Orientation-sensitive interval width and computation
Ahmadnasab and Psarrakos [10], 2024	Location results for special structured pencils under linear perturbation	Demonstrates structure-dependent spectral bounds	General definite-Hermitian low-rank formulation
Gernandt and Trunk [11], 2025	Parametric rank-one motion and interlacing	Closest recent source for low-rank interlacing behavior	Simultaneous matrix-pair updates of arbitrary low support dimension
Hochstenbach et al. [12], 2025	Structure-preserving methods for singular pencils	Emphasizes preservation under Hermitian perturbations	Definite regular pair and certified selected-eigenvalue intervals
Dodig and Stošić [13], 2025	Minimal-rank pencil perturbations	Extends algebraic perturbation/completion theory	Numerical localization after update data are known
Demmel et al. [14], 2026	Randomized inverse-free pencil diagonalization	Represents modern full-pencil algorithmic direction	Cheaper repeated localization after baseline spectral preprocessing

It is seen that the present contribution is an intersection result in the comparison. It is a combination of definiteness of pair min–max argument, low rank projection factors, rank interlacing and small inertia correction. None of these ingredients by itself is the sole contribution, it is the novelty of the combination of these into an ordered enclosure and refinement framework.

3. Objectives, Scope, and Mathematical Formulation

3.1 Research objectives

The major goal is to establish and verify the strong framework of enclosure of definite ordered Hermitian matrix pencils for structured low rank perturbations. Specific Objectives include:

- O1. Make sure that the definiteness for a true ordered spectrum has been maintained and tested.
- O2. Compute two-sided eigenvalue intervals which are valid on the global level as well as for projections, for simultaneous perturbations of both matrices.
- O3. Low support dimension: exploit getting a determinant and/or inertia formula with reduced dimension to count eigenvalues.
- O4. Narrow (perturbed) intervals without performing the full (perturbed) eigendecomposition.
- O5. Generalization of positive definite denominator formulation to arbitrary definiteness “Hermitian pairs”.
- O6. Assess coverage and interval width, clustered eigenvalue effect, physical interpretability and cost of repeated update.

3.2 Canonical positive-definite setting

The central analysis begins with

$$Ax = \lambda Bx, \quad A = A^*, \quad B = B^* > 0. \quad (3)$$

The ordered eigenvalues are denoted by

$$\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n. \quad (4)$$

There exists a B -orthonormal eigenvector matrix X such that

$$AX = BX\Lambda, \quad X^*BX = I, \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_n). \quad (5)$$

The structured perturbations are represented by

$$E = UC_U U^*, \quad F = VD_V V^*, \quad (6)$$

where $C_U = C_U^*$, $D_V = D_V^*$, and $r_A = \text{rank}(E)$ and $r_B = \text{rank}(F)$ are much smaller than n . Whitening by $B^{-1/2}$ gives

$$H = B^{-1/2}AB^{-1/2} = Q\Lambda Q^*, \quad (7)$$

where $Q = B^{1/2}X$ is unitary. Let W be an orthonormal basis for the union of the ranges of $B^{-1/2}EB^{-1/2}$ and $B^{-1/2}FB^{-1/2}$. Then, for $s \leq r_A + r_B$,

$$B^{-1/2}EB^{-1/2} = WCW^*, \quad B^{-1/2}FB^{-1/2} = WDW^*, \quad (8)$$

with $C = C^*$ and $D = D^*$ of order s . The perturbed problem is therefore equivalent to

$$(H + WCW^*)y = \tilde{\lambda}(I + WDW^*)y. \quad (9)$$

This is a common support representation, which is NOT a limitation: separate update ranges are combined using a thin QR factorization and compressed into C and D .

3.3 Exact and sufficient definiteness conditions

Because $W^*W = I_s$, the denominator update acts as the identity on $\mathcal{R}(W)^\perp$ and as $I_s + D$ on $\mathcal{R}(W)$. Consequently,

$$I + WDW^* > 0 \quad \Leftrightarrow \quad I_s + D > 0. \quad (10)$$

A convenient sufficient condition is

$$\delta := \|D\|_2 < 1. \quad (11)$$

This reduced test addresses O1 and avoids checking an $n \times n$ updated denominator.

3.4 Enclosure objective and evaluation measures

For the ordered perturbed eigenvalues

$$\tilde{\lambda}_1 \leq \tilde{\lambda}_2 \leq \dots \leq \tilde{\lambda}_n, \quad (12)$$

the objective is to compute

$$\ell_i \leq \tilde{\lambda}_i \leq u_i, \quad i = 1, \dots, n. \quad (13)$$

For an interval $I_i = [\ell_i, u_i]$, its width is $w_i = u_i - \ell_i$. Relative performance is measured against the shifted global interval $I_i^{(g)}$ by

$$\text{width ratio}_i = \frac{w_i}{w_i^{(g)}}, \quad \text{improvement}_i = \frac{w_i^{(g)}}{w_i}. \quad (14)$$

Coverage is reported as the fraction of directly computed reference eigenvalues that lie in the proposed intervals, with a small tolerance proportional to machine precision. The direct eigenvalues are used only for experimental validation and not by the enclosure algorithm.

4. Main Theoretical Results

4.1 Shifted global enclosure

Choose a real shift $\tau < \lambda_1$ and define

$$H_\tau = H - \tau I > 0, \quad C_\tau = C - \tau D, \quad (15)$$

and

$$\rho_\tau = \|H_\tau^{-1/2}WC_\tau W^*H_\tau^{-1/2}\|_2. \quad (16)$$

Theorem 1 (shifted global enclosure). Assume $\delta < 1$ and $\rho_\tau < 1$. Then, for every $i = 1, \dots, n$,

$$\tau + \frac{1-\rho_\tau}{1+\delta}(\lambda_i - \tau) \leq \tilde{\lambda}_i \leq \tau + \frac{1+\rho_\tau}{1-\delta}(\lambda_i - \tau). \quad (17)$$

Proof. The definition of ρ_τ implies the Loewner inequalities

$$(1 - \rho_\tau)H_\tau \preceq H_\tau + WC_\tau W^* \preceq (1 + \rho_\tau)H_\tau. \quad (18)$$

Similarly, $\|D\|_2 = \delta$ and $W^*W = I_s$ give

$$(1 - \delta)I \leq I + WDW^* \leq (1 + \delta)I. \quad (19)$$

Applying the generalized Courant–Fischer principle to the shifted perturbed quotient yields

$$\frac{1-\rho_\tau}{1+\delta}(\lambda_i - \tau) \leq \tilde{\lambda}_i - \tau \leq \frac{1+\rho_\tau}{1-\delta}(\lambda_i - \tau), \quad (20)$$

which proves the result.

Theorem 1 provides a baseline that respects the relative scale induced by the shift. It is global because the same two factors ρ_τ and δ control all indices.

4.2 Projection-aware enclosure

Let $q_1, \dots, q_n$ be the columns of Q , and define the lower and upper invariant-subspace matrices

$$Q_{1:i} = [q_1, \dots, q_i], \quad Q_{i:n} = [q_i, \dots, q_n]. \quad (21)$$

The update-space overlap factors are

$$a_i = \|W^* Q_{1:i}\|_2^2, \quad b_i = \|W^* Q_{i:n}\|_2^2. \quad (22)$$

Both quantities lie in $[0,1]$. They are small when the low-rank support is nearly orthogonal to the corresponding trial subspace. Let

$$\mu_i = \lambda_i - \tau, \quad \gamma_\tau = \|C_\tau\|_2. \quad (23)$$

Theorem 2 (projection-aware two-sided enclosure). Assume $\delta < 1$, $1 - \delta a_i > 0$, and $\mu_i - \gamma_\tau b_i > 0$. Then

$$\tau + \frac{\mu_i - \gamma_\tau b_i}{1 + \delta b_i} \leq \tilde{\lambda}_i \leq \tau + \frac{\mu_i + \gamma_\tau a_i}{1 - \delta a_i}. \quad (24)$$

Proof. For a unit vector $y \in \mathcal{R}(Q_{1:i})$, one has

$$y^* H_\tau y \leq \mu_i, \quad \|W^* y\|_2^2 \leq a_i. \quad (25)$$

Therefore,

$$y^* (H_\tau + W C_\tau W^*) y \leq \mu_i + \gamma_\tau a_i \quad (26)$$

and

$$y^* (I + WDW^*) y \geq 1 - \delta a_i. \quad (27)$$

The min–max characterization, using $\mathcal{R}(Q_{1:i})$ as an admissible i -dimensional trial subspace, gives the upper bound. For a unit vector $y \in \mathcal{R}(Q_{i:n})$,

$$y^* H_\tau y \geq \mu_i, \quad \|W^* y\|_2^2 \leq b_i. \quad (28)$$

Thus the numerator is at least $\mu_i - \gamma_\tau b_i > 0$, while the denominator is at most $1 + \delta b_i$. The max–min characterization on the $(n - i + 1)$ -dimensional upper invariant subspace yields the lower bound. $\square$

The theorem explicitly quantifies the advantage of low-rank orientation. A perturbation can have a non-negligible global norm but a small effect on one side of an eigenvalue interval when its range has weak overlap with the corresponding invariant trial space.

Computational remark. Let

$$M = Q^* W \in \mathbb{C}^{n \times s}. \quad (29)$$

Then a_i and b_i are the largest eigenvalues of the cumulative $s \times s$ Gram matrices

$$\sum_{j=1}^i m_j^* m_j \quad \text{and} \quad \sum_{j=i}^n m_j^* m_j, \quad (30)$$

where m_j is row j of M . Hence all overlap factors are computed in $O(ns^2 + ns^3)$ operations with $O(ns)$ storage, and no dense $n \times n$ projection is required.

4.3 Low-rank interlacing guard

The generalized eigenvalues of the updated pair are the ordinary eigenvalues of

$$T = (I + WDW^*)^{-1/2} (H + W C W^*) (I + WDW^*)^{-1/2}. \quad (31)$$

Theorem 3 (rank guard). If $I + WDW^* > 0$, then

$$\text{rank}(T - H) \leq 2s. \quad (32)$$

Consequently, with $k = \min(2s, n - 1)$,

$$\lambda_{i-k} \leq \tilde{\lambda}_i \leq \lambda_{i+k}, \quad (33)$$

where a missing lower or upper index is interpreted as $-\infty$ or $+\infty$, respectively.

Proof. Functional calculus on the invariant decomposition $\mathcal{R}(W) \oplus \mathcal{R}(W)^\perp$ gives

$$(I + WDW^*)^{-1/2} = I + WSW^* \quad (34)$$

for some Hermitian matrix S of order s . Expanding $T - H$ shows that its range is contained in $\mathcal{R}(W) + \mathcal{R}(HW)$, whose dimension is at most $2s$. The rank inequality follows. Standard Hermitian rank- k interlacing applied to T and H yields the eigenvalue guard. $\square$

The rank guard is independent of perturbation magnitude. It is especially useful near spectral edges and when the analytic intervals become wide.

4.4 Combined a priori enclosure

Let $[\ell_i^{(g)}, u_i^{(g)}]$ denote the interval from Theorem 1, $[\ell_i^{(p)}, u_i^{(p)}]$ the interval from Theorem 2, and $[\ell_i^{(r)}, u_i^{(r)}]$ the rank guard from Theorem 3. Their intersection gives

$$\ell_i = \max\{\ell_i^{(g)}, \ell_i^{(p)}, \ell_i^{(r)}\}, \quad u_i = \min\{u_i^{(g)}, u_i^{(p)}, u_i^{(r)}\}. \quad (35)$$

Whenever the hypotheses of Theorems 1 and 2 hold, $[\ell_i, u_i]$ is a valid enclosure and cannot be wider than any constituent interval. This intersection is called the **combined enclosure** below.

4.5 Reduced determinant and inertia formulas

The a priori intervals can be refined using the low-rank secular structure. For a real test point t satisfying $t \notin \{\lambda_1, \dots, \lambda_n\}$, define

$$K(t) = C - tD, \quad (36)$$

and assume temporarily that $K(t)$ is nonsingular. With $M = Q^*W$, define

$$R(t) = M^*(\Lambda - tI)^{-1}M, \quad S(t) = K(t)^{-1} + R(t). \quad (37)$$

Theorem 4 (compressed determinant and inertia correction). At every test point for which $H - tI$ and $K(t)$ are nonsingular,

$$\det(H - tI + WK(t)W^*) = \det(H - tI)\det K(t)\det S(t), \quad (38)$$

and

$$J(H - tI + WK(t)W^*) = J(H - tI) + J(-S(t)) - J(-K(t)^{-1}). \quad (39)$$

Here $J(G) = (p(G), q(G), z(G))$ denotes the Hermitian inertia, where p , q , and z count positive, negative, and zero eigenvalues, respectively.

Proof. The determinant identity follows from the matrix determinant lemma. For inertia, consider the Hermitian block matrix

$$\mathcal{H}(t) = \begin{bmatrix} H - tI & W \\ W^* & -K(t)^{-1} \end{bmatrix}. \quad (40)$$

Taking the Schur complement first with respect to $H - tI$ gives the complementary block $-S(t)$, whereas taking it with respect to $-K(t)^{-1}$ gives $H - tI + WK(t)W^*$. Equating the two Haynsworth inertia decompositions yields the stated correction formula. $\square$

Because the updated denominator is positive definite, the number of perturbed eigenvalues strictly below t equals the negative index q of the updated shifted matrix. At a counting point that is not a perturbed eigenvalue, $S(t)$ is nonsingular by the determinant identity; since $K(t)$ is also nonsingular, the negative-index component of the inertia formula may be written using $p(G) = s - q(G)$. Therefore, if

$$N(t) = |\{j: \tilde{\lambda}_j < t\}|, \quad (41)$$

then Theorem 4 gives the exact count

$$N(t) = N_0(t) + q(K(t)^{-1}) - q(S(t)), \quad (42)$$

where $N_0(t) = |\{j: \lambda_j < t\}|$. The only update-dependent inertia calculations are of order s . Moreover,

$$R(t) = \sum_{j=1}^n \frac{m_j^* m_j}{\lambda_j - t}, \quad (43)$$

so one count costs $O(ns^2 + s^3)$ operations.

Test points coinciding with an original eigenvalue or a root of $\det K(t)$ are avoided by a one-sided perturbation of the test point or by splitting the interval at the pole. This does not alter the limiting eigenvalue count.

4.6 Count-certified interval refinement

Theorem 5 (reduced-inertia bisection). Suppose an interval $[a, b]$ satisfies

$$N(a) < i \leq N(b) \quad (44)$$

for a target index $i \in \{1, \dots, n\}$. Repeatedly set $m = (a + b)/2$ and replace a by m when $N(m) < i$, or replace b by m otherwise. After j iterations,

$$\tilde{\lambda}_i \in [a_j, b_j], \quad b_j - a_j = 2^{-j}(b - a) \quad (45)$$

in exact arithmetic.

Proof. The function $N(t)$ is nondecreasing and increases by the algebraic multiplicity when t crosses an eigenvalue. The count condition is therefore preserved by the stated update, and ordinary bisection halves the width at every step. $\square$

The initial interval can be the combined enclosure. In floating-point arithmetic, the implementation uses Hermitian LDL^* or eigenvalue-based inertia for the small matrices, moves test points away from poles, and rounds final endpoints outward. A formally verified implementation would additionally require interval arithmetic; the experiments in Section 6 use double precision and validate every reported enclosure against an independent dense generalized eigensolve.

4.7 Extension to a general definite Hermitian pair

A Hermitian pair (A, B) is definite when a real rotation exists such that one rotated coefficient is positive definite. Define

$$A_\theta = A \cos \theta - B \sin \theta, \quad B_\theta = A \sin \theta + B \cos \theta. \quad (46)$$

Choose θ such that $B_\theta > 0$. If λ is an eigenvalue of (A, B) and μ is the corresponding eigenvalue of (A_θ, B_θ) , then

$$\mu = \frac{\lambda \cos \theta - \sin \theta}{\lambda \sin \theta + \cos \theta}, \quad \lambda = \frac{\sin \theta + \mu \cos \theta}{\cos \theta - \mu \sin \theta}. \quad (47)$$

The perturbations rotate as

$$E_\theta = E \cos \theta - F \sin \theta, \quad F_\theta = E \sin \theta + F \cos \theta. \quad (48)$$

Their combined range has dimension at most $r_A + r_B$. On any interval that does not cross the pole $\cos \theta - \mu \sin \theta = 0$, the inverse Möbius map is strictly increasing because

$$\frac{d\lambda}{d\mu} = \frac{1}{(\cos \theta - \mu \sin \theta)^2} > 0. \quad (49)$$

Thus a valid interval for μ_i maps endpointwise to a valid interval for λ_i . This establishes O5 without changing the low-rank order of the update.

5. Computational Framework

5.1 Projection-aware low-rank pencil enclosure algorithm

The resulting workflow is called **Projection-Aware Low-Rank Pencil Enclosure (PALPE)**. It separates one-time preprocessing from update-dependent work.

Algorithm 1: PALPE for selected perturbed eigenvalues

Input: A definite Hermitian pair (A, B) ; baseline ordered eigenpairs; low-rank factors of E and F ; target indices J ;

tolerance ε .

1. If necessary, rotate the pair so that the denominator is positive definite.
2. Whiten the pair and form $H = Q\Lambda Q^*$ during baseline preprocessing.
3. Form an orthonormal basis W for the union of the whitened update ranges, and compress the updates to C and D .
4. Test $I_s + D > 0$. Choose $\tau < \lambda_1$ and verify $\rho_\tau < 1$.
5. Compute $M = Q^*W$, the overlap sequences a_i, b_i , the shifted global intervals, the projection-aware intervals, and the rank guards.
6. Intersect the three intervals to obtain $[\ell_i, u_i]$.
7. For each requested $i \in \mathcal{I}$, evaluate the reduced count $N(t)$ at interval endpoints. Expand slightly only if finite-precision endpoint equality prevents the count condition.
8. Apply reduced-inertia bisection until $u_i - \ell_i \leq \varepsilon \max(1, |\ell_i|, |u_i|)$.
9. Map the endpoints back through the inverse Möbius transformation when a rotation was used.

Output: Ordered intervals containing the selected perturbed eigenvalues under the exact-arithmetic hypotheses, together with diagnostic values δ , ρ_τ , overlap factors, and iteration counts.

5.2 Complexity

A complete baseline generalized eigendecomposition costs $O(n^3)$ dense operations and is assumed to be available once. For each new low-rank update, constructing the common support and the compressed matrices costs $O(ns^2 + s^3)$ after whitening operations are available. Computing $M = Q^*W$ costs $O(n^2s)$ for a dense explicitly stored Q , or less when eigenvectors are applied implicitly or only local spectral information is retained. Once M is formed, all overlap sequences cost $O(ns^2 + ns^3)$ with small constants because the eigenproblems have order s .

Each reduced count requires $O(ns^2 + s^3)$ operations, dominated by the rational sum defining $R(t)$. For p target eigenvalues and j bisection steps, the update-dependent refinement cost is

$$O(pj(ns^2 + s^3)). \quad (50)$$

This is to compare with $O(n^3)$ for a full dense recomputation. The lowrank procedure is thus the most appealing if s is much less than n , if the baseband spectral information remains the same, and the small number of requested eigenvalues is desired. A highly tuned dense eigensolver can be faster for small n which is measured in Section 6.5 and not assumed.

5.3 Numerical stability safeguards

Five safeguards are used in the implementation: (i) A reduced positive-definiteness test for $I_s + D$; (ii) A shift that is separate from λ_1 ; (iii) Explicit check of the criteria in Theorem 2, $\rho_\tau < 1$ and lower-numerator criteria; (iv) One-sided movement away from poles of $R(t)$ or $K(t)^{-1}$; and (v) Rounding of the outermost endpoints outward. The inertia is computed symmetrically as it is the case for $S(t)$ which is Hermitian. Experiments are performed in double precision (no interval-arithmetic verification is made).

6. Numerical Results

6.1 Experimental design

Experiments were all repeated with IEEE double precision and fixed random seeds. The Monte Carlo run was done for both real symmetric and complex Hermitian cases. Indeed, a dense Hermitian generalized eigensolver was used to compute the reference perturbed eigenvalues independently. Tests were done on a single computational thread in order to minimize variations.

Baseline spectra were ordered and the low-rank update cores were indefinite Hermitian matrices. A number of denominator perturbations were scaled to keep $I_s + D$ positive definite. An evaluation of the entire enclosure was performed for each index. The following measures have been taken:

1. coverage of the direct reference eigenvalues;
2. mean and median width reduction relative to Theorem 1;

3. effect of support dimension and orientation;
4. count agreement on independent test grids;
5. final width and iterations for clustered eigenvalues; and
6. per-update time for eight selected eigenvalues compared with a full dense eigensolve.

All these experiments do O2–O6. They do not simply provide a visual example to test, but are also written to test the consequences of the theorem.

6.2 Effect of update-space orientation

Problem instances were created with support dimension $s=3$ in 120×120 instances where the support is aligned mainly below a certain target value, or even near the target, or above or even randomly. Average widths and improvement factors are provided in Table 2. There was no intervillage or image contacting problem found.

Support orientation	Mean global width	Mean projection width	Mean combined width	Median improvement	90th-percentile improvement	Failures
Lower spectral region	0.503650	0.226557	0.208417	2.261	2.889	0
Target region	0.448341	0.227695	0.203020	2.018	2.563	0
Upper spectral region	0.418760	0.230155	0.199510	2.018	2.374	0
Random support	0.638202	0.272319	0.263272	2.141	3.186	0




Figure 1. Width profiles for a representative random support. The projection-aware and rank-intersected intervals are substantially narrower than the shifted global enclosure over most of the spectrum.

The table works with this mechanism, as explained in figure 1. The width in the global scale increases with the spectral scale since it is determined by the same relative factor. The projection-aware width is still quite appealing, as a_i and b_i model the extent of the visible lower and upper half of the update support. This finding directly contributes to O2.

6.3 Rank-and-magnitude study

For the systematic study: $n=96$, four support dimensions, three update magnitudes and 18 pencils per configuration. 50% of the trials were Complex Hermitian trials. Altogether 216 pencils were tested; there were 20,736 ordered perturbed eigenvalues. The results are summarised in Table 3.

s	Core scale	Eigenvalues tested	Coverage (%)	Mean reduction (%)	Median improvement	90th-percentile improvement
1	0.05	1728	100.0	56.92	2.502	3.270
1	0.10	1728	100.0	56.65	2.470	3.209
1	0.15	1728	100.0	57.29	2.528	3.292
3	0.05	1728	100.0	51.30	2.119	2.996
3	0.10	1728	100.0	49.71	2.026	2.841
3	0.15	1728	100.0	50.56	2.089	2.905
5	0.05	1728	100.0	47.64	1.955	2.802
5	0.10	1728	100.0	46.30	1.890	2.685
5	0.15	1728	100.0	47.05	1.913	2.763
8	0.05	1728	100.0	44.00	1.815	2.623
8	0.10	1728	100.0	43.95	1.805	2.678
8	0.15	1728	100.0	43.77	1.784	2.651

The average reduction was 49.60% for all cases. There was no change in the median improvement across the range of 1.784 to 2.528. The reduction slows down with increasing s because more of the trial subspaces overlap (with increasing s , the update space overlaps increasing amounts of both the trial subspaces). The behavior of the relative improvement with the core scale was only slightly different since both the global and the projection-induced widths increase with the magnitude of the perturbation.




Figure 2. Median global-to-combined width ratio versus low-rank support dimension. Orientation information is most valuable for the smallest ranks, but remains beneficial throughout the tested range.

The zero failure count in the empirical data agrees with Theorems 1–3, and the systematic trend in the ranks offers an interpretation that can be handed by using the global norm by itself.

6.4 Clustered eigenvalues and reduced-inertia refinement

A 140×140 test pencil was constructed with six eigenvalues separated by only 2×10^{-6} around the value 8. The combined a priori intervals necessarily overlapped because the perturbation was larger than the unperturbed cluster gaps. The reduced count was then used to isolate each ordered perturbed eigenvalue. Table 4 reports the result.

Index	Perturbed eigenvalue	Initial width	Refined width	Iterations	Contained
68	7.996573924868	2.2723e-1	5.2909e-11	32	Yes
69	7.999998766842	2.2191e-1	5.1671e-11	32	Yes
70	8.000000431834	2.2204e-1	5.1701e-11	32	Yes
71	8.000003119670	2.2227e-1	5.1756e-11	32	Yes
72	8.000231544553	2.2199e-1	5.1690e-11	32	Yes
73	8.002973891176	2.2147e-1	5.1568e-11	32	Yes

The mean refinement factor was approximately 4.29×10^9 . On 401 independent grid points, excluding poles, the reduced count agreed with the direct reference count at every tested point. The result shows that interval overlap is not a failure of the a priori theory: the exact count can separate ordered eigenvalues without assuming a large spectral gap.




Figure 3. Initial combined widths and reduced-inertia refined widths for the six-eigenvalue cluster.

This experiment directly addresses O3 and O4 and complements the cluster verification emphasis of [6] with a low-rank update-specific mechanism.

6.5 Localized mass-spring perturbation

A fixed-end chain with $n = 180$ degrees of freedom and nonuniform positive masses was used as a physically interpretable generalized eigenproblem. A rank-one spring modification and a rank-one local mass addition produced updates to both matrices. After whitening, their combined support dimension was two. Table 5 reports six modes distributed across the spectrum.

Mode	Original eigenvalue	Perturbed eigenvalue	Exact shift	Global width	Combined width	Reduction (%)	Refined width
5	0.0067181314	0.0067174953	-6.3605e-7	0.155647	0.014234	90.85	6.63e-12
20	0.1063428196	0.1064534353	1.1062e-4	0.182741	0.064695	64.60	3.77e-12
50	0.6300449700	0.6300917596	4.6790e-5	0.325168	0.167822	48.39	4.88e-12
90	1.7645616015	1.7654651960	9.0359e-4	0.633712	0.246190	61.15	7.17e-12
140	3.0952365623	3.0982786518	3.0421e-3	0.995604	0.138151	86.12	1.61e-11
175	3.8493712715	3.8348896490	-1.4482e-2	1.200699	0.149236	87.57	1.74e-11

The mean reduction for the reported modes was 73.11%, and every reduced interval contained the independently computed perturbed eigenvalue. The reduced count also agreed with the reference count on all 300 tested non-pole points.




Figure 4. Combined/global width ratio over all vibration modes. Localized physical changes interact very differently with different modal subspaces.

This figure illustrates the reason for sometime using a single global norm in localized models. Much narrower intervals are given for modes whose pattern of their displacement doesn't have much to do with the modified spring or mass direction. This will directly relate the mathematical overlap factors to the physical interpretation of this update and to O6.

6.6 Repeated-update timing

The comparison of the timing experiment was a full dense generalized eigensolve (which returns all n eigenvalues) and PALPE on 8 target indices following baseline spectral pre-processing. Median single thread times are listed in Table 6. In all cases, the errors resulting from the enclosure tolerance were smaller than 2.3×10^{-9} .

n	Targets	Full eigensolve (s)	PALPE workflow (s)	Full/PALPE ratio	Mean count evaluations
200	8	0.00342	0.02112	0.162	216.0
400	8	0.01995	0.02596	0.769	215.8
600	8	0.05993	0.03105	1.930	215.3

n	Targets	Full eigensolve (s)	PALPE workflow (s)	Full/PALPE ratio	Mean count evaluations
800	8	0.13238	0.03821	3.464	215.5




Figure 5. Median time per update. A highly optimized dense eigensolver is faster at small orders, while the target-only low-rank workflow crosses over as the matrix order increases.

The result prevents an overgeneralized efficiency claim. At $n = 200$ and $n = 400$, the dense solver was faster because of optimized level-3 linear algebra and the fixed overhead of about 216 reduced counts. At $n = 600$ and $n = 800$, the target-only workflow was approximately 1.93 and 3.46 times faster, respectively. The benefit is expected to increase when updates are repeated, s remains small, and only a few eigenvalues are required. A sparse large-scale implementation would require different baselines and is left for future work.

7. Discussion

7.1 Interpretation in relation to the research objectives

O1 was obtained with the same reduced condition above, that is, $I_s + D > 0$, and the convenient condition that successively used, which is $\|D\|_2 < 1$. This is more informative than testing the updated $n \times n$ denominator directly and maintains the ordered real spectrum which is essential in the min-max theory.

O2 Theorems 1-3 discussed issues related to . The tested cases demonstrate that the numerical evidence reveals that the projection-aware and rank-intersected bounds always bound the global bound, and that the projection-aware and rank intersected bounds are generally narrower. Better gains are found at smaller dimensions of support which is consistent with the overlap mechanism. However, the results of the experiments indicate that orientation is also important regardless of the size of perturbation.

O3 and O4 were achieved by Theorems 4 and 5. The clustered example was the most explicit case with the $s \times s$ inertia correction isolating all ordered eigenvalues with initial gaps being much larger than the gaps in the clusters. The method does not involve the first order separation assumption.

O5 The rotation and the Möbius map were used to achieve an analytical solution for . Although the $B > 0$ canonical case is the one more often used in numerical experiments, the general definite pair is reduced with no loss of low-rank property, as shown in the derivation.

O6 Which was addressed via complementary experiments: real and complex random pencils tested for broad validity; orientation and rank study accounted for the behavior of width; difficult spectral configuration study via the cluster;

physical meaning of the phenomenon provided by the mass–spring study; and timing study identified both the benefit and the crossover limitation.

Objective	Theoretical mechanism	Numerical evidence	Main conclusion
O1: preserve definiteness	Reduced test $I_S + D \succ 0$	All generated denominators passed the reduced test	Definiteness checking is low dimensional
O2: sharper ordered intervals	Theorems 1–3 and combined intersection	20,736 eigenvalues; 49.6% mean width reduction	Low-rank orientation materially reduces conservatism
O3: reduced eigenvalue count	Theorem 4	Zero count mismatch on cluster and mechanical grids	Updated counts can be recovered from $S \times S$ inertia corrections
O4: target refinement	Theorem 5	Cluster widths near 5.2×10^{-11} after 32 steps	Overlapping a priori intervals can be separated reliably
O5: general definite pair	Rotation and monotone Möbius map	Analytical extension	Positive-denominator theory covers general definite pairs by transformation
O6: practical relevance	PALPE workflow	Physical model and timing crossover	Benefit is greatest for repeated low-rank, target-only updates

7.2 Relation to previous studies

The results presented here complement the literature on algebraic perturbation. Some studies like [2,7-9,13,15-18] analyze the possible occurrence of any eigenvalues and/or eigenvectors or invariant factors under certain rank restriction. PALPE takes as an input an update, the update is assumed to be already known, and PALPE asks for the location of the ordered real eigenvalues. It thus provides quantitative information, which is not aimed in those canonical-structure results.

The major difference with the definite-pair perturbation studies [1],[5] is that, in this case, Q^*W matrices and their cumulative Gram matrices are explicitly used. As the same-norm update is rotated w.r.t the eigenbasis, the interval for the index i will change, but will not be apparent in a purely global norm. Compared to the verified full-eigenpair approach of [6] the method proposed here employs an update of low dimension and precomputed baseline spectral data, leaving the task of selected-eigenvalue refinement as the core. Thus, PALPE can be considered as an update/certification layer, and not as a standalone universal eigensolver, as it exists for the diagonalization of dense complex matrices [12] or modern eigensolvers for full-pencil diagonalization [14].

The rank guard is similar to recent focus on interlacing for rank one pencil perturbations [11] but extends the use of interlacing in computation to a support of dimension s and simultaneous rank perturbations of both the coefficient. The reduced determinant and inertia formula provides the missing link between low rank structure and ordered interval refinement.

7.3 Limitations

There are five boundaries to this current contribution.

First, the most efficient use is based on an eigendecomposition or similar spectral representation of baseline. If matrices are the only available data, and just one update is needed, it may be better to calculate the eigensolve all the way. Second, Theorems 1 and 2 need to be specified sufficient inequalities. However, it is convenient to think of an adaptive shift selection, as large changes in the denominators or an inappropriate choice of the shift may be used to violate them even though the new pair is still definite. Third, the exact reduced count formula does not give test points at which $K(t)$ is singular. The splitting handles of Safeguarded splitting are only for handling isolated poles; a production implementation should use a pivoted formulation where the term $K(t)$ is not explicitly inverted. Outward rounding and independent validation, rather than complete interval arithmetic, is used for the finite-precision code. Last, there are a lot of experiments packed in a small space. For sparse, large scale problems, we would need to introduce here local eigenspaces, or rational approximation of the sum of the resolvent, or hierarchical versions of evaluating $R(t)$.

These restrictions are not restrictions on the mathematical validity (under the stated hypotheses) but concern of the choice of the appropriate computational regime.

8. Conclusions

In this paper they came up with a unified approach for eigenvalue enclosures of definite Hermitian matrix pencils of structured low-rank perturbations. A scale aware baseline was provided by a shifted global bound. A new projection-aware bound was developed, which included index dependent overlap factors between the update support and invariant sub-spaces below and above. A movement-control rank interlacing guard additionally further curtailed the movement allowed. This led us to a workable a priori interval that is the intersection of their intervals.

A compressed determinant-inertia identity was applied in the second part of the framework to be able to reconstruct the number of the perturbed eigenvalues from the original ordered spectrum together with an $s \times s$ correction. Selected eigenvalues were then further bisected to the desired accuracy, yielding results that were compared with known values. The approach was further developed in the case of the positive definite denominator by replacing it by a real rotation and a monotone function of the Mehbius type.

All the numerical results obtained agreed with the theory. The results indicated no coverage failure across 20,736 randomly selected real and complex Hermitian eigenvalues, the global widths were reduced by 49.6% on average over these randomly-selected examples, when the combined intervals were used. Eigenvalue clustering was done to a separation of about 5×10^{-11} ; localized mass-spring modifications were seen to have significant gains as a function of mode number. The timing study also provided an understanding of regime: the low-rank approach is not necessarily faster on small matrices; however, when only a range of eigenvalues is desired, and a number of repeated updates are made to the resulting matrices, the low rank method provides an advantage in timing.

It would be interesting to develop interval-arithmetic inertia certification, adaptive shift optimization, pole-free reduced factorizations, partial-eigenspace variants, and sparse implementations, in which a factorization of the compressed resolvent is done without saving an eigenvector matrix.

References

- [1] C. Xiao and H. Li, “New perturbation bounds of partitioned generalized Hermitian eigenvalue problem,” *Applied Mathematics and Computation*, vol. 321, pp. 422–430, 2018, doi: 10.1016/j.amc.2017.10.063.
- [2] I. Baragaña and A. Roca, “Weierstrass structure and eigenvalue placement of regular matrix pencils under low rank perturbations,” *SIAM Journal on Matrix Analysis and Applications*, vol. 40, no. 2, pp. 440–453, 2019, doi: 10.1137/18M1200245.
- [3] F. De Terán, C. Mehl, and V. Mehrmann, “Low rank perturbation of regular matrix pencils with symmetry structures,” *Foundations of Computational Mathematics*, vol. 22, no. 1, pp. 257–311, 2022, doi: 10.1007/s10208-021-09500-4.
- [4] C. Mehl, V. Mehrmann, and M. Wojtylak, “Matrix pencils with coefficients that have positive semidefinite Hermitian parts,” *SIAM Journal on Matrix Analysis and Applications*, vol. 43, no. 3, pp. 1186–1212, 2022, doi: 10.1137/21M1439997.
- [5] I. Kuzmanović Ivičić, S. Miodragović, and M. Ugrić, “The $\tan\theta$ theorem for definite matrix pairs,” *Linear and Multilinear Algebra*, vol. 71, no. 15, pp. 2484–2498, 2023, doi: 10.1080/03081087.2022.2106349.
- [6] S. M. Rump and M. Lange, “Fast computation of error bounds for all eigenpairs of a Hermitian and all singular pairs of a rectangular matrix with emphasis on eigen- and singular value clusters,” *Journal of Computational and Applied Mathematics*, vol. 434, Art. no. 115332, 2023, doi: 10.1016/j.cam.2023.115332.
- [7] M. Dodig and M. Stošić, “Bounded rank perturbations of quasi-regular pencils over arbitrary fields,” *SIAM Journal on Matrix Analysis and Applications*, vol. 44, no. 4, pp. 1879–1907, 2023, doi: 10.1137/22M1504068.
- [8] F. De Terán, A. Dmytryshyn, and F. M. Dopico, “Generic eigenstructures of Hermitian pencils,” *SIAM Journal on Matrix Analysis and Applications*, vol. 45, no. 1, pp. 260–283, 2024, doi: 10.1137/22M1523297.
- [9] M. Dodig and M. Stošić, “Bounded rank perturbations of a regular matrix pencil,” *Linear Algebra and its Applications*, vol. 689, pp. 126–154, 2024, doi: 10.1016/j.laa.2024.02.028.
- [10] M. Ahmadnasab and P. J. Psarrakos, “Eigenvalue characterization of some structured matrix pencils under linear perturbation,” *Electronic Journal of Linear Algebra*, vol. 40, pp. 274–298, 2024, doi: 10.13001/ela.2024.7371.
- [11] H. Gernandt and C. Trunk, “Eigenvalues of parametric rank-one perturbations of matrix pencils,” *Linear Algebra and its Applications*, vol. 708, pp. 429–457, 2025, doi: 10.1016/j.laa.2024.12.012.
- [12] M. E. Hochstenbach, C. Mehl, and B. Plestenjak, “Solving singular generalized eigenvalue problems. Part III: Structure preservation,” *SIAM Journal on Matrix Analysis and Applications*, vol. 46, no. 3, pp. 1964–1987, 2025, doi: 10.1137/24M1668603.
- [13] M. Dodig and M. Stošić, “Minimal rank perturbations of matrix pencils,” *Linear Algebra and its Applications*, vol. 721, pp. 81–101, 2025, doi: 10.1016/j.laa.2025.03.002.

- [14] J. Demmel, I. Dumitriu, and R. Schneider, “Generalized pseudospectral shattering and inverse-free matrix pencil diagonalization,” *Foundations of Computational Mathematics*, vol. 26, pp. 569–645, 2026, doi: 10.1007/s10208-024-09682-7.
- [15] I. Baragaña and A. Roca, “Fixed rank perturbations of regular matrix pencils,” *Linear Algebra and its Applications*, vol. 589, pp. 201–221, 2020, doi: 10.1016/j.laa.2019.12.022.
- [16] M. Dodig and M. Stošić, “Rank one perturbations of matrix pencils,” *SIAM Journal on Matrix Analysis and Applications*, vol. 41, no. 4, pp. 1889–1911, 2020, doi: 10.1137/19M1279411.
- [17] I. Baragaña, M. Dodig, A. Roca, and M. Stošić, “Bounded rank perturbations of regular pencils over arbitrary fields,” *Linear Algebra and its Applications*, vol. 601, pp. 180–188, 2020, doi: 10.1016/j.laa.2020.05.015.
- [18] S. Moon and S. Park, “Lower and upper bounds for the number of distinct eigenvalues of a perturbed regular matrix pencil,” *Linear Algebra and its Applications*, vol. 626, pp. 171–187, 2021, doi: 10.1016/j.laa.2021.05.017.
- [19] A. C. M. Ran and M. Wojtylak, “Global properties of eigenvalues of parametric rank one perturbations for unstructured and structured matrices II,” *Complex Analysis and Operator Theory*, vol. 16, no. 6, Art. no. 91, 2022, doi: 10.1007/s11785-022-01268-x.