Off-diagonal perturbation, first-order approximation and quadratic residual bounds for matrix eigenvalue problems

Yuji Nakatsukasa

Abstract When a symmetric block diagonal matrix $\left[\begin{smallmatrix} A_1 & \\ & A_2 \end{smallmatrix}\right]$ undergoes an off-diagonal perturbation $\left[\begin{smallmatrix} A_1 & E_{12} \\ E_{21} & A_2 \end{smallmatrix}\right]$, the eigenvalues of these matrices are known to differ only by $O(\frac{\|E_{12}\|^2}{\text{gap}})$, which scales *quadratically* with the norm of the perturbation. Here gap measures the distance between eigenvalues, and plays a key role in the constant. Closely related is the first-order perturbation expansion for simple eigenvalues of a matrix. It turns out that the accuracy of the first-order approximation is also $O(\frac{\|E\|^2}{\text{gap}})$, where E is the perturbation matrix. Also connected is the residual bounds of approximate eigenvalues obtained by the Rayleigh-Ritz process, whose accuracy again scales quadratically in the residual, and inverse-proportionally with the gap between eigenvalues. All these are tightly linked, but the connection appears to be rarely discussed. This work elucidates this connection by showing that all these results can be understood in a unifying manner via the quadratic perturbation bounds of block diagonal matrices undergoing off-diagonal perturbation. These results are essentially known for a wide range of eigenvalue problems: symmetric eigenproblems (for which the explicit constant can be derived), nonsymmetric and generalized eigenvalue problems. We also extend such results to matrix polynomials, and show that the accuracy of a first-order expansion also scales as $O(\frac{\|E\|^2}{\text{gap}})$, and argue that two-sided projection methods are to be preferred to one-sided projection for nonsymmetric eigenproblems, to obtain higher accuracy in the computed eigenvalues.

keywords: quadratic eigenvalue perturbation bounds, off-diagonal perturbation, first-order expansion, eigenvalue gap, polynomial eigenvalue problems

Yuji Nakatsukasa
Department of Mathematical Informatics, The University of Tokyo, Tokyo 113-8656, Japan
e-mail: nakatsukasa@mist.i.u-tokyo.ac.jp

1 Introduction

Classical eigenvalue perturbation theory studies bounds or approximations to the eigenvalues and eigenvectors of $A+E$ for some "small" E (such as small norm or low rank), given the knowledge of some information on A, such as an eigenpair such that $Ax_i = \lambda_i(A)x_i$. Such results are of interest in large-scale scientific computing because, for example, (i) given the information of A, they give estimates for the eigenvalues and eigenvectors of $A+E$ that can be obtained cheaply (for example the first-order expansion (1)), and (ii) they can be used to give quantitative analysis for the accuracy of the computed eigenpairs. See [7] and [14, Ch. IV] for an overview of eigenvalue perturbation theory.

The eigenvalues of two unstructured matrices A and $A+E$ generally differ by $O(\|E\|_2)$, or sometimes more (as large as $O(\|E\|_2^{1/n})$ in the worst case, when defective eigenvalues are present). However, there are important situations when eigenvalues behave more nicely than such general bounds suggest, and this work focuses on such cases.

Among the most well-known results for the (simplest and most well-understood) symmetric case $A = A^T \in \mathbb{R}^{n\times n}$ are Weyl's theorem $|\lambda_i(A) - \lambda_i(A+E)| \le \|E\|$ (throughout, we employ the spectral norm $\|A\| = \sigma_{\max}(A)$ for matrices, and 2-norm for vectors), and the first-order perturbation expansion for simple eigenvalues (e.g. [3, §7.2.2])

$$\lambda_i(A+E) = \lambda_i(A) + \frac{x^T E x}{x^T x} + O(\|E\|^2). \tag{1}$$

Here x is a nonzero eigenvector such that $Ax = \lambda_i(A)x$. Note that this gives an approximation $\lambda_i(A) + \frac{x^T A x}{x^T x}$ to $\lambda_i(A+E)$, rather than a bound as in Weyl's theorem.

This work revolves around the less well-known (but, we argue, equally important) quadratic perturbation bounds for eigenvalues of block-diagonal matrices that undergo off-diagonal perturbation.

Theorem 1 ([6, 9]). *Let $A, E \in \mathbb{R}^{n\times n}$ are symmetric matrices partitioned as*

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix}, \quad E = \begin{bmatrix} 0 & E_1^T \\ E_1 & 0 \end{bmatrix}, \tag{2}$$

then [9]

$$|\lambda_i(A+E) - \lambda_i(A)| \le \frac{\|E\|^2}{\mathrm{gap}_i}, \tag{3}$$

and a slightly tighet bound [6] holds:

$$|\lambda_i(A+E) - \lambda_i(A)| \le \frac{2\|E\|^2}{\mathrm{gap}_i + \sqrt{\mathrm{gap}_i^2 + 4\|E\|^2}} \le \frac{\|E\|\ 2}{\mathrm{gap}_i}. \tag{4}$$

Here gap_i *is defined by*

$$\mathrm{gap}_i := \begin{cases} \min_{\lambda_j \in \lambda(A_2)} |\lambda_i - \lambda_j| & \text{if} \quad \lambda_i \in \lambda(A_1) \\ \min_{\lambda_j \in \lambda(A_1)} |\lambda_i - \lambda_j| & \text{if} \quad \lambda_i \in \lambda(A_2). \end{cases} \tag{5}$$

Note that the bound in (3) scales quadratically with the perturbation $\|E\|$, which is significantly smaller than Weyl's bound $\|E\|$ when $\|E\|$ is small. We shall look at (3) from many different viewpoints, and one central goal of this work is to reveal the implications of such quadratic bounds. For simplicity, we refer to situations as in (2), where a block-diagonal matrix undergoes off-diagonal perturbation, as *off-diagonal perturbation*.

We note that the bounds in (4) are sharp in the sense that without further information, there are examples where equality is attained. This sharpness can be confirmed by verifying that the first bound in (4) is exact for 2×2 matrices, and it reduces to the second bound in the limit $\|E\| \to 0$.

Also of interest in this work is quadratic residual bounds [5], which claims for symmetric eigenvalue problems that defining the residual by $r = A\hat{x} - \hat{\lambda}\hat{x}$ where $\hat{x}$ is an approximate eigenvector with $\|\hat{x}\| = 1$ and $\hat{\lambda} = \hat{x}^T A\hat{x}$ is the Rayleigh quotient, we have

$$|\lambda - \hat{\lambda}| \leq \frac{\|r\|^2}{\mathrm{gap}_i}, \tag{6}$$

where gap_i is as in (5) with A_1 taken to be 1×1. The notion of gap is subtly different between (6) and those in the literature, e.g. [1, § 4.8],[12, Thm. 11.7.1]: we explain this more in Section 2.1. More generally, with the Ritz values $\{\hat{\lambda}_i\}_{i=1}^k$ for a symmetric matrix obtained as the eigenvalues of $X^T AX$ where $X \in \mathbb{R}^{n\times k}$ has orthonormal columns, we have

$$|\lambda - \hat{\lambda}| \leq \frac{\|R\|^2}{\mathrm{gap}_i}, \tag{7}$$

where gap_i is again as in (5), which is "widened", resulting in an improved bound. We derive this below in Section 2.2.

The first-order expansion (1) and the off-diagonal quadratic perturbation bounds (3) are closely connected: specifically, the first-order perturbation expansion (1) explains why off-diagonal perturbation results in quadratic eigenvalue perturbation bounds (without information on the constant $\frac{1}{\mathrm{gap}_i}$). Conversely, using (3) one can obtain (1), and moreover obtain the constant $\frac{1}{\mathrm{gap}_i}$ hidden in the trailing term $O(\|E\|^2)$. We shall see that the residual bound can also be regarded as a consequence of (3). In other words, (3) can be regarded as a fundamental fact that implies many results in eigenvalue perturbation theory.

These connections are known to experts in eigenvalue perturbation theory, but to the author's knowledge there is no literature that states them explicitly. One goal of this work is to clarify this connection, which holds not only for symmetric eigenvalue problems but also for nonsymmetric and generalized eigenvalue problems. All this is not exactly new, in that they are simply observations that connect results in the literature.

The second goal of this work is to extend such results to polynomial eigenvalue problems. For polynomial eigenvalue problems, the first-order perturbation expan-

sion (1) is known [15], but no result seems to be available on off-diagonal perturbation analogous to (3). We shall obtain such result, and show that if

$$P(\lambda) = \begin{bmatrix} P_1(\lambda) & 0 \\ 0 & P_2(\lambda) \end{bmatrix}, \quad E = \begin{bmatrix} 0 & E_{12}(\lambda) \\ E_{21}(\lambda) & 0 \end{bmatrix}, \tag{8}$$

then

$$\lambda_i(P) - \lambda_i(P+E) \le c\frac{\|E(\lambda_i(P))\|^2}{\text{gap}_i}, \tag{9}$$

for some c, which depends on the conditioning of the eigenvectors. The point we wish to convey here is that the eigenvalue gap plays the same role even in polynomial eigenvalue problems. Note that $E(\lambda_i(P))$ is the value of the matrix polynomial $E(\lambda)$ (representing the perturbation) evaluated at $\lambda = \lambda_i(P)$.

All in all, in this note we investigate the quadratic eigenvalue perturbation bounds under off-diagonal perturbation such as (3) and (9) from different viewpoints, and reveal some of their practical ramifications.

The only reason we stated the above results for symmetric matrices is for simplicity; extensions to nonsymmetric and generalized eigenproblems are available [8, 9, 14]. We structure this note similarly: we first discuss the symmetric case in Section 2, then deal with nonsymmetric and generalized eigenvalue problems, and treat polynomial eigenvalue problems at the end.

In what follows, for simplicity we normalize any right eigenvector to have unit norm, and we scale the left eigenvector via the orthogonality relation such as $y^T x = 1$ or $y^T P(\lambda_i)' x = 1$. $\lambda_i(A)$ denotes the ith eigenvalue of A, arranged in ascending order if A is symmetric, and otherwise its ordering is insignificant: $\lambda_i(A)$ denotes a specific eigenvalue of A.

2 Symmetric case

We start by treating the simplest case of symmetric eigenvalue problems; entirely analogous results hold for the complex Hermitian case.

As advertised, let us first explain how (1) implies off-diagonal perturbation should result in quadratic eigenvalue perturbation bounds. Let the matrices A, E be as in (2), and $\lambda_i(A) \in \lambda(A_1)$ with $\text{gap}_i > 0$. Then $Ax = \lambda_i(A)x$ with eigenvector structure $x = \begin{bmatrix} x_1 \\ 0 \end{bmatrix}$, and so by substituting into (1) we obtain

$$\begin{aligned} \lambda_i(A+E) &= \lambda_i(A) + x^T E x + O(\|E\|^2) \\ &= \lambda_i(A) + \begin{bmatrix} x_1 \\ 0 \end{bmatrix}^T \begin{bmatrix} 0 & E_1^T \\ E_1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ 0 \end{bmatrix} + O(\|E\|^2) \\ &= \lambda_i(A) + O(\|E\|^2), \end{aligned}$$

in which we note that $\begin{bmatrix} x_1 \\ 0 \end{bmatrix}^T \begin{bmatrix} 0 & E_1^T \\ E_1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ 0 \end{bmatrix} = 0$ due to the block structure. That is, the first-order term in (1) disappears because of the structure. Thus the first term in the perturbation expansion scales quadratically with the perturbation $\|E\|$.

2.1 First-order expansion and its constant via off-diagonal bounds

We now turn to the connection in the opposite direction and derive the first-order expansion (1) using the quadratic bounds (3). We are not claiming the derivation here is simpler than the standard method of differentiating the equation $Ax = \lambda x$ and left-multiplying the left eigenvector (see [3, § 7.2.2] or [15]). However, as we shall see, the derivation here reveals the constant in front of the quadratic term $\|E\|^2$. In Section 3.4 we also give an explanation based on Gerschgorin's theorem.

Using (3) we shall derive the following result, which can be seen as a variant of (1) that reveals the constant hidden in $O(\|E\|^2)$. Note below that $\widetilde{\text{gap}}_i$ can be regard as a modified gap.

Proposition 1. *Let $A, E \in \mathbb{R}^{n\times n}$ be symmetric matrices. Then*

$$|\lambda_i(A+E) - (\lambda_i(A) + \frac{x^T E x}{x^T x})| \leq \frac{\|E\|^2}{\widetilde{\text{gap}}_i}, \tag{10}$$

where $\widetilde{\text{gap}}_i = \max(0, \min_{j\neq i} |\lambda_i + \frac{x^T E x}{x^T x} - \lambda_j| - \|E\|)$, and $\widetilde{\text{gap}}_i \to \text{gap}_i$ as $E \to 0$.

Proof. Consider the eigenvalue decomposition

$$X^T A X = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix},$$

where $X = [x_1, \ldots, x_n]$ is an orthogonal eigenvector matrix. Then

$$X^T (A+E) X = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} + X^T E X,$$

whose (i,i) element is $\lambda_i + x_i^T E x_i$. We can then apply a permutation matrix P that moves the ith position to the first (the specific choice of P does not matter), which gives

$$P^T X^T (A+E) X P = \begin{bmatrix} \lambda_i + x_i^T E x_i & & x_i^T E x_j \\ & \ddots & \\ x_i^T E x_j & & \lambda_n + x_n^T E x_n \end{bmatrix}.$$

we now partition this matrix and write $P^TX^T(A+E)XP = \begin{bmatrix} A_1 & \\ & A_2 \end{bmatrix} + \begin{bmatrix} 0 & E_1^T \\ E_1 & 0 \end{bmatrix}$, where $A_1 = \lambda_i + x_i^T E x_i$ is 1×1 (highlighted in red), a scalar, hence $A_1 = \lambda_i + x_i^T E x_i$. Noting that $\lambda_i(P^TX^T(A+E)XP) = \lambda_i(A+E)$, we now use Theorem 1 to obtain

$$|\lambda_i(A+E) - (\lambda_i(A) + x_i^T E x_i)| \le \frac{\|E\|^2}{\widetilde{\mathrm{gap}}_i},$$

where $\widetilde{\mathrm{gap}}_i := \max(0, \mathrm{gap}_i - \|E\|)$. This updated gap is obtained by using Weyl's bound for the lower-right $(n+1)\times(n+1)$ part of $P^TX^T(A+E)XP$, which is altered from A_2 by the lower-right part of P^TX^TEXP. This establishes (18) (and hence also the first-order expansion (1)). □

Note that $\widetilde{\mathrm{gap}}_i$ is different from gap_i: as alluded to after (6), this difference is reflected in the formal statements of the residual bounds and quadratic off-diagonal perturbation bounds in the following sense: in (6) the gap is between an approximate and exact eigenvalue. In (5) the gap is between two approximate eigenvalues. While this subtlety is certainly present, we shall not expound on this further as this difference diminishes as $E \to 0$. Furthermore, they both convey the same message that the accuracy scaled inverse-proportionally to the gap.

2.2 Connection to residual bounds

Now we explain how the residual bounds (6), (7) can be obtained from the off-diagonal quadratic perturbation bound (3).

Recall that the Rayleigh-Ritz process employs a subspace spanned by a matrix $Q \in \mathbb{R}^{n\times k}$ with orthonormal columns with $k < n$, and computes the $k\times k$ symmetric eigenproblem $Q^TAQ = V_Q\Lambda_Q V_Q^T$, from which one extracts the Ritz values $\mathrm{diag}(\Lambda_Q)$, and Ritz vectors QV_Q. Here we examine the accuracy of the Ritz values, and derive (6) using (3).

Consider $Q^\perp \in \mathbb{R}^{n\times(n-k)}$, which spans the orthogonal complement of Q so that $[Q\ Q^\perp]$ is a square orthogonal matrix. Then we have

$$[Q\ Q^\perp]^T A [Q\ Q^\perp] = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}.$$

Then $Q^TAQ = A_{11}$, and the problem essentially reduces to quantifying the accuracy of the eigenvalues of A_{11} as approximants to k of the eigenvalues of A. This is exactly the problem treated in Theorem 1, and gives (6): again, the gap in the literature differs subtly from what we get here. The above argument more generally gives (7). Note how the definition of gap differs between (6) and (7); the gap is usually much wider in (7), giving better bounds if the residual norms are of comparable magnitudes.

3 Non-symmetric eigenvalue problems

The main message of the last section was the significance of the off-diagonal quadratic perturbation bound (3) in symmetric eigenvalue problems. We now turn to the analogous results for more general eigenvalue problems, focusing on nonsymmetric standard eigenproblems.

3.1 Statements

Here we display the extensions of results in the previous section to nonsymmetric matrices. When A is nonsymmetric, the first-order perturbation expansion (1) becomes

$$\lambda_i(A+E) = \lambda_i(A) + \frac{y^T E x}{y^T x} + O(\|E\|^2), \tag{11}$$

where $Ax = \lambda_i(A)x$ as before and y is a left eigenvector, that is, $y^T A = y^T \lambda_i(A)$. Here and below, $\lambda_i(A)$ denotes an eigenvalue of A (not necessarily ordered, as they can be nonreal) and $\lambda_i(A+E)$ denotes an eigenvalue that spawns from $\lambda_i(A)$ in that $\lambda_i(A+tE)$ is continuous in $t \in \mathbb{R}$ and $\lambda_i(A+tE) \to \lambda_i(A)$ as $t \to 0$. The expansion (11) holds for any scaling of x, y; we scale them so that $\|x\| = 1$ and $y^T x = 1$.

The analogue of the key result (3) becomes [3, Ch. 7] the following (ignoring the nontrivial issue of "ordering" the eigenvalues, which can be nonreal). Note below that neither A nor E is assumed to be symmetric, but the block structure is preserved.

Theorem 2 ([9]). *Let $A, E \in \mathbb{R}^{n\times n}$ be matrices partitioned as*

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix}, \quad E = \begin{bmatrix} 0 & E_{12} \\ E_{21} & 0 \end{bmatrix}. \tag{12}$$

Then [9, Thm. 5]

$$\lambda_i(A+E) - \lambda_i(A) \le c\frac{\|E_{12}\|\|E_{21}\|}{\mathrm{gap}_i}, \tag{13}$$

where c is the product of condition numbers of eigenvector matrices of $A+E$ and A.

Theorem 13 is lower in sophistication than Theorem 1 in terms of sharpness: for example, it is not exact for 2×2 case. It nonetheless suffices for the argument here.

Bounds like (13) are often stated in terms of the quantity sep [13, § 4.2], which here is $\mathrm{sep}_i = 1/\|(A_2 - \lambda_i(A))^{-1}\|$. Note that $\mathrm{sep}_i = \mathrm{gap}_i$ when the matrices are symmetric (or normal), and sep_i takes into account the conditioning of the eigenvector matrix. In this note, for simplicity we absorb this effect in the constant c, in order to highlight the role played by the gap throughout eigenvalue problems.

Finally, the residual bound in the nonsymmetric case becomes [13, Thm. 4.2.12]

$$|\lambda - \hat{\lambda}| \le c\frac{\|r_r\|\|r_l\|}{\mathrm{gap}_i}, \tag{14}$$

where

$$\hat{\lambda} = \frac{y^T A x}{y^T x}$$

is the Ritz value (via two-sided projection) and r_r, r_l are the right and left residual vectors defined by

$$r_r = Ax - \hat{\lambda} x, \qquad r_l = y^T A - \hat{\lambda} y^T. \tag{15}$$

More generally, for block matrices (or projection onto a $k > 1$-dimensional subspace is employed), we have

$$|\lambda_i - \hat{\lambda}_i| \le c \frac{\|R_r\| \|R_l\|}{\mathrm{gap}_i}, \tag{16}$$

where

$$\hat{\lambda}_i = \lambda_i(Y^T A X - \lambda Y^T X)$$

are the eigenvalues of the matrix pencil $Y^T AX - \lambda Y^T X$ (sometimes called Ritz values via two-sided projection), and denoting $\hat{\Lambda} = \mathrm{diag}(\hat{\lambda}_1, \ldots, \hat{\lambda}_k)$, R_r, R_l are the right and left residual matrices defined by

$$R_r = AX - X\hat{\Lambda}, \qquad R_l = Y^T A - \hat{\Lambda} Y^T. \tag{17}$$

Below we follow the same line of argument as in Section 2 and derive the first-order expansion (11) and residual bound (14) using the off-diagonal quadratic perturbation bound (13).

3.2 First-order expansion and its constant via off-diagonal bounds

Let us establish an analogue of Proposition for the nonsymmetric case.

Proposition 2. *Let $A, E \in \mathbb{R}^{n\times n}$ be symmetric matrices. Then*

$$|\lambda_i(A+E) - (\lambda_i(A) + \frac{x^T E x}{x^T x})| \le \frac{\|E\|^2}{\widetilde{\mathrm{gap}}_i}, \tag{18}$$

where $\widetilde{\mathrm{gap}}_i = \max(0, \min_{j\neq i} |\lambda_i + \frac{x^T E x}{x^T x} - \lambda_j| - \|E\|)$, and $\widetilde{\mathrm{gap}}_i \to \mathrm{gap}_i$ as $E \to 0$.

To establish (11), assume that A is diagonalizable (this assumption is mainly for simplicity: it can be relaxed to just $\lambda_i(A)$ being simple, or even to multiple eigenvalues as long as they are not defective) and consider the eigenvalue decomposition

$$X^{-1} A X = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix},$$

where $X = [x_1, \ldots, x_n]$ is a nonsingular eigenvector matrix. Then

$$X^{-1}(A+E)X = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} + X^{-1}EX,$$

where recalling that $[y_1,\ldots,y_n]^T = X^{-1}$, the (i,i) element is $\lambda_i + y_i^T E x_i$. We can then apply a permutation matrix P that moves the ith position to the first, which gives

$$P^T X^{-1}(A+E)XP = \begin{bmatrix} \lambda_i + y_i^T E x_i & & y_i^T E x_j \\ & \ddots & \\ y_i^T E x_j & & \lambda_n + y_n^T E x_n \end{bmatrix}. \tag{19}$$

We now partition this matrix and write $P^T X^T (A+E)XP = \begin{bmatrix} A_1 & \\ & A_2 \end{bmatrix} + \begin{bmatrix} 0 & E_1^T \\ E_1 & 0 \end{bmatrix}$, where A_1 is 1×1, a scalar, hence $A_1 = \lambda_i + y_i^T E x_i$. Noting that $\lambda_i(P^T X^{-1}(A+E)XP) = \lambda_i(A+E)$, we now use Theorem 2 to obtain

$$|\lambda_i(A+E) - (\lambda_i(A) + y_i^T E x_i)| \le c\frac{\|E\|^2}{\widetilde{\text{gap}}_i},$$

where $\widetilde{\text{gap}}_i := \max(0, \text{gap}_i - \tilde{c}\|E\|)$; this is a lower bound for the gap between $\lambda_i(A) + y_i^T E x_i)$ and the eigenvalues of the $(n-1)\times(n-1)$ bottom-right part of (19), and $\tilde{c}$ depends on its eigenvector matrix. This establishes (18), and hence also the first-order expansion (1).

3.3 Connection to residual bounds

Now we explain how the residual bound (14) can be obtained from (13).

For the nonsymmetric case, we analyze the two-sided projection method, which spanned by two matrices: $X \in \mathbb{R}^{n\times k}$, hoped to approximate some *right* eigenvectors), usually but not necessarily with orthonormal columns, and $Y \in \mathbb{R}^{n\times k}$ (hoped to approximate the same *left* eigenvectors; however, the simple choice $Y = X$ is quite common and natural in view of the Schur form.

We then compute the $k\times k$ generalized eigendecomposition $V_Y^T(Y^TAX, Y^TX)V_X = (\Lambda_{XY}, I)$, which reduces to a standard eigenproblem if we choose Y so that $Y^TX = I$. One then extracts the approximate eigenvalues (sometimes also called Ritz values) as $\text{diag}(\Lambda_{XY})$, and approximate right and left eigenvectors (Ritz vectors) XV_X and YV_Y. Here we examine the accuracy of the Ritz values, and derive (16) using Theorem 2.

For simplicity we discuss the case $Y^TX = I_k$. Let X_2, Y_2 be such that $[X\ X_2], [Y\ Y_2]$ are nonsingular matrices and $[Y\ Y_2]^T[X\ X_2] = I_n$. Then we write

$$[Y\ Y_2]^T A[X\ X_2] = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}.$$

Then $Y^T AX = A_{11}$, and the problem essentially reduces to quantifying the accuracy of the eigenvalues of A_{11} as approximants to k of the eigenvalues of A. This is exactly the problem treated in Theorem 2, in which $\|A_{21}\|$ corresponds to the right residual $\|R_r\|$ and $\|A_{21}\|$ to $\|R_l\|$, leading to (16).

We note that the residual bounds become linear in $\|R_r\|$ if we use a one-sided projection method with $Y = X$, as then $\|R_l\|$ will be $O(1)$ rather than $O(\|R_r\|)$. This indicates that it is worth using two-sided projection when an approximation to the left eigenvectors is available.

3.4 Gerschgorin's viewpoint

Here we explain the same quadratic scaling $|\lambda_i(A+E) - \widehat{\lambda}_i| \le \frac{\|E\|^2}{\mathrm{gap}_i}$ from the viewpoint of Gerschgorin's theorem. We could have included such treatment in the symmetric case, but we have deferred its treatment until now since no simplification accrues in the symmetric case. Gerschgorin's theorem states that

$$\lambda(A) \in \bigcup_i \Gamma_i, \qquad \Gamma_i = \{z \in \mathbb{C} |\ |z - a_{ii}| \le \sum_{j \ne i} |a_{ij}|\},$$

that is, the eigenvalues of A lie in the union of Gerschgorin disks Γ_i of radius $\sum_{j\ne i}|a_{ij}|$ centered at a_{ii}. Now we focus on λ_i, and denoting by ε an entry bounded by $|\varepsilon| \le \|E\|$, we see that

$$P^T X^{-1}(A+E)XP = \begin{bmatrix} \lambda_i & & & \\ & \lambda_1 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix} + \begin{bmatrix} \varepsilon & \varepsilon & \cdots & \varepsilon \\ \varepsilon & \varepsilon & \cdots & \varepsilon \\ \vdots & \vdots & \ddots & \vdots \\ \varepsilon & \varepsilon & \cdots & \varepsilon \end{bmatrix}.$$

If λ_i is a simple eigenvalue and E is sufficiently small, we will have $\Gamma_j \cap \Gamma_i = \phi$ for $j \ne i$, which means there is exactly one eigenvalue lying in Γ_i. Let δ be a quantity smaller than $\mathrm{gap}_i = \min_{j\ne i}|\lambda_i - \lambda_j|$. Then using the diagonal matrix $D = \mathrm{diag}(\frac{\varepsilon}{\delta}, 1, \ldots, 1)$ we have

$$DP^T X^{-1}(A+E)XPD^{-1} = \begin{bmatrix} \lambda_i & & & \\ & \lambda_1 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix} + \begin{bmatrix} \varepsilon_{ii} & \frac{\varepsilon^2}{\delta} & \cdots & \frac{\varepsilon^2}{\delta} \\ \delta & \varepsilon & \cdots & \varepsilon \\ \vdots & \vdots & \ddots & \vdots \\ \delta & \varepsilon & \cdots & \varepsilon \end{bmatrix}.$$

Now with Gerschgorin's theorem applied to this matrix, $\Gamma_i \cap \Gamma_j = \phi$ still holds, and $\Gamma_i : |z-(\lambda_i+\varepsilon_{ii})| \leq O(\frac{\varepsilon^2}{\delta})$. Note the radius of Γ_i is now $O(\frac{\varepsilon^2}{\mathrm{gap}_i})$. Noting that $\varepsilon_{ii} = y^T Ex/(y^T x)$ where x, y are the left/right eigenvectors of A corresponding to λ_i, it follows that $\lambda_i + \varepsilon_{ii} = \lambda_i + y^T Ex/(y^T x)$ approximates an eigenvalue of $A+E$ to $O(\frac{\varepsilon^2}{\mathrm{gap}_i})$ accuracy.

The above diagonal scaling technique combined with Gerschgorin's theorem is again commonly used, for example in [14, Ch. IV].

3.5 Extensions to generalized eigenproblem

Analogous results for generalized eigenvalue problems can be established, using quadratic off-diagonal perturbation bounds presented in [8]. In particular, the Gerschgorin argument can be used for establishing quadratic perturbation bounds for generalized nonsymmetric eigenvalue problems; see the last section of [10]. We omit the details here.

4 Polynomial eigenvalue problems

We now turn to polynomial eigenvalue problems. In a polynomial eigenvalue problem, one is to find $x \neq 0$ such that $P(\lambda)x = 0$ where $P(\lambda) = \sum_{i=0}^{k} \lambda^i A_i \in \mathbb{C}[\lambda]^{n\times n}$ is a matrix polynomial. Let $E(\lambda) \in \mathbb{C}[\lambda]^{n\times n}$ be another matrix polynomial, representing a perturbation to $P(\lambda)$. The first-order perturbation expansion of an eigenvalue $\lambda_i(P(\lambda))$, with $\lambda_i((P+tE)(\lambda))$ depending continuously on t (as in (11)), is known [15] to be

$$\lambda_i((P+E)(\lambda)) = \lambda_i(P(\lambda)) - \frac{y^T E(\lambda_i P(\lambda))x}{y^T P'(\lambda)x} + O(\|E(\lambda_i P(\lambda))\|_2^2). \tag{20}$$

The denominator $y^T P'(\lambda)x$ in the first-order term is known to be nonzero when $\lambda_i(P(\lambda))$ is simple. The expansion (20) is in fact valid without the restriction that $E(\lambda)$ is a matrix polynomial of the same or less degree as $P(\lambda)$, but here we focus on such cases (as otherwise the number of eigenvalues is not controlled). We can verify that (20) reduces to the expansions (1) and (11) in the special case where $P(\lambda)$ represents a linear standard eigenvalue problem $P(\lambda) = \lambda I - A$.

4.1 Analysis via linearization

The most common approach to studying polynomial eigenvalue problems, both in theory and practice is linearization [2]. Here we follow this standard approach to examine the accuracy of first-order expansion (20) and to derive quadratic perturbation bounds for matrix polynomials. The most well known and widely used linearization is the companion linearization. For a *monic* matrix polynomial $P(\lambda) = \sum_{i=0}^{k} A_i \lambda^i$ with $A_k = I$, the companion linearization is defined by

$$C = \begin{bmatrix} -A_{k-1} & -A_2 & -A_1 & -A_0 \\ I & & & \\ & \ddots & & \\ & & I & \end{bmatrix}. \tag{21}$$

This $kn \times kn$ matrix clearly has kn eigenvalues, which match those of $P(\lambda)$, so we can write $\lambda_i(C) = \lambda_i(P(\lambda))$. The right eigenvectors of $P(\lambda)$ and C are related by the Vandermonde structure as follows: if $P(\lambda_i)x_i = 0$, then

$$C \begin{bmatrix} \lambda_i^{k-1} x \\ \vdots \\ \lambda_i x \\ x \end{bmatrix} = \lambda_i \begin{bmatrix} \lambda_i^{k-1} x \\ \vdots \\ \lambda_i x \\ x \end{bmatrix}. \tag{22}$$

In view of the first-order expansion, we also need the left eigenvector of C. Let y be a left eigenvector of P such that $y^T P(\lambda) = 0$. Then the left eigenvector of C has the structure of the Horner shift [4, eq. (3.12)]

$$\begin{bmatrix} y \\ (\lambda_i A_k + A_{k-1})y \\ \vdots \\ (\lambda_i^k A_k + \lambda^{k-1} A_{k-1} + \cdots + A_1)y \end{bmatrix}^T C = \lambda_i \begin{bmatrix} y \\ (\lambda_i A_k + A_{k-1})y \\ \vdots \\ (\lambda_i^k A_k + \lambda^{k-1} A_{k-1} + \cdots + A_1)y \end{bmatrix}^T. \tag{23}$$

We denote the right and left eigenvectors of C by $\underline{x}$ and $\underline{y}$ respectively, and use (22) and (23) for (20) to obtain the first-order expansion of the eigenvalue λ_i of P as

$$\lambda_i((P+E)(\lambda)) = \lambda_i + \frac{y^T(\sum_{j=0}^{k} \lambda_i^j E_i)x}{y^T(\sum_{j=1}^{k} \lambda_i^{i-1} A_i)x} + O(\|E(\lambda_i)\|^2). \tag{24}$$

On the other hand, denoting by $C + \Delta C$ the companion linearization associated with $P + E$, the expansion with respect to C becomes (using (11) with $A \leftarrow C$)

$$\lambda_i(C + \Delta C) = \lambda_i + \frac{\underline{y}^T(\Delta C)\underline{x}}{\underline{y}^T\underline{x}} + O(\|\Delta C\|^2), \tag{25}$$

which, in view of (22) and (23), is equivalent to (24); this is to be expected because $P+E$ and $C+\Delta C$ have the same eigenvalues.

The value in the equivalence between (24) and (25) is that with (25), we can invoke the analysis for linear eigenvalue problems to examine the eigenvalues of P and its perturbed variant. Indeed, assuming λ_i is a simple eigenvalue, the exact same arguments as in Section 3 shows that the second-order term in the expansion (25) can be written as $O(c\frac{\|\Delta C\|^2}{\text{gap}_i})$. Note that this allows for general perturbation in the matrix C, whereas the perturbation of interest here is structured, because, as we can see in (23), the only elements in C that depend on P are those in the first block row. In any case, we have proven the following result.

Theorem 3. *Let $P(\lambda)=\sum_{i=0}^{k}\lambda^i A_i \in \mathbb{C}[\lambda]^{n\times n}$ be a monic matrix polynomials of degree k, and $E(\lambda)=\sum_{i=0}^{k-1}\lambda^i E_i \in \mathbb{C}[\lambda]^{n\times n}$. Let (λ_i,x_i,y_i) be a simple eigentriple of $P(\lambda)$. Then*

$$\lambda_i((P+E)(\lambda)) = \lambda_i + \frac{y^T(\sum_{j=0}^{k}\lambda_i^j E_i)x}{y^T(\sum_{j=1}^{k}\lambda_i^{j-1}A_i)x} + O(c\frac{\|E(\lambda_i)\|^2}{\text{gap}_i}), \tag{26}$$

where $\text{gap}_i = \min_{j\neq i}|\lambda_i-\lambda_j(P(\lambda))|$ and c depends on the conditioning of the eigenvector matrix of C in (21)*.*

We have not yet examined whether the un-structured perturbation results in a constant that is smaller than the unstructured counterpart by (25) would indicate. To examine whether this happens, we turn to MATLAB experiments in which we construct a random matrix polynomial $P(\lambda)$ companion matrix as in (21), compute an eigentriple (λ_i,x_i,y_i), then examine the perturbation in λ_i when we introduce perturbation in C in two different forms:

1. Perturb only the first block row by norm ε,
2. Perturb the whole matrix C by norm ε,

for some small ε, which here we set to 10^{-4}. We then examine the difference in the accuracy of $\lambda_i + y^T(\Delta C)x/(y^Tx)$ as an approximation to an eigenvalue of the perturbed matrix; clearly, since the second type includes the first as a special case, the second would lead to a larger perturbation in the worst case. We experimented with various n and $k=2,3,\ldots,10$, with randomly generated perturbation matrices, and observed that there is never a significant difference between the sensitivity of λ_i under the two types of perturbations. That said, making this observation precise seems nontrivial, and we leave it as an open problem.

4.2 Quadratic bounds by off-diagonal perturbation

We now turn to off-diagonal perturbation and derive a bound analogous to (9).

Theorem 4. *Let*

$$P(\lambda)=\begin{bmatrix} P_1(\lambda) & 0 \\ 0 & P_2(\lambda)\end{bmatrix}, \quad E(\lambda)=\begin{bmatrix} 0 & E_{12}(\lambda) \\ E_{21}(\lambda) & 0\end{bmatrix} \tag{27}$$

be matrix polynomials, with $P(\lambda)$ being monic and degree k, and $E(\lambda)$ of degree $k-1$ or less. Then

$$\lambda_i(P(\lambda))-\lambda_i((P+E)(\lambda)) \le c\frac{\|E(\lambda_i(P))\|^2}{\mathrm{gap}_i}, \tag{28}$$

where $\mathrm{gap}_i=\min_{j\neq i}|\lambda_i(P(\lambda))-\lambda_j(P(\lambda))|$ *and c depends on the conditioning of the eigenvector matrix of C in* (21)*.*

Proof. The argument is simple as we now have all the essential tools. Note that for any eigenvalue of $P_1(\lambda)$ that is not an eigenvalue of $P_2(\lambda)$, the left and right eigenvectors have the block zero structure $\begin{bmatrix} x_1 \\ 0\end{bmatrix}$ and $\begin{bmatrix} y_1 \\ 0\end{bmatrix}$. Plugging this into (26), we obtain (28). □

Note that the above argument takes the opposite route from before: now we are using the first-order expansion to obtain the quadratic off-diagonal perturbation bound. We conjecture that the gap_i can be replaced with the "widened" gap as in Theorem 1. However, it appears to be nontrivial to obtain a direct proof of (28) (with the refined gap) by extending the argument in [6, 9], which rely on the Sylvester law of intertia.

Observe in (28) that what matters for the perturbation in λ_i is the magnitude of $E(\lambda)$ evaluated at $\lambda=\lambda_i$; for example, the perturbation is zero if $E(\lambda_i)=0$, even if $E(\lambda)$ takes large values away from λ_i.

4.2.1 Accuracy of eigenvalues obtained by projection methods

Another implication of Theorem 4 can be observed on an approximate eigenpair $(\hat{\lambda}_i,\hat{x}_i)$ obtained via a projection method applied to polynomial eigenvalue problems. Consider for simplicity a symmetric matrix polynomial $P(\lambda)$. Suppose $(\hat{\lambda}_i,\hat{x}_i)$ is obtained by solving $V^TP(\hat{\lambda}_i)Vy_i=0$ for some orthonormal matrix $V\in\mathbb{C}^{n\times k}$, $k<n$, with $\hat{x}_i=Vy_i$. Then we can write, using an orthogonal matrix $[V\ V^\perp]$,

$$[V\ V^\perp]^TP(\lambda)[V\ V^\perp]=\begin{bmatrix} P_1(\lambda) & E_{12}(\lambda) \\ E_{21}(\lambda) & P_2(\lambda)\end{bmatrix},$$

where $P_1(\hat{\lambda}_i)$ has $\hat{\lambda}_i$ as an exact eigenvalue, and the residual $\|P(\lambda_i)\hat{x}_i\|$ (which is computable) is bounded by $\|E_{12}(\hat{\lambda}_i)\|=\|E_{21}(\hat{\lambda}_i)\|$ (usually not computable). Thus by the above theorem it follows that the computed eigenvalue $\hat{\lambda}_i$ has accuracy $O(\frac{\|E(\lambda_i(P))\|^2}{\mathrm{gap}_i})=O(\frac{\|P(\hat{\lambda}_i)\hat{x}_i\|^2}{\mathrm{gap}_i})$.

Note that the same type of quadratic bound follows for *nonsymmetric* matrix polynomials, provided that we employ a two-sided projection method in which

we work with $Y^T P(\lambda)X$ where Y and X approximate the desired left and right eigenspaces respectively. This is exactly the same situation as in linear eigenvalue problems, for which we need two-sided projection to obtain quadratic eigenvalue convergence in the nonsymmetric case. Put another way, because the left and right eigenvectors are the same for symmetric eigenvalue problems, the Rayleigh-Ritz method automatically approximates both the left and right eigenvectors simultaneously. The apparent difference in convergence speed for symmetric and nonsymmetric eigenvalue problems (which is present e.g. in the QR algorithm and Rayleigh quotient iteration) comes from the fact that the algorithm is implicitly employing a one-sided projection method, not because the convergence is inherently hindered by lack of symmetry.

5 Discussion

This work examined the ramifications of the fact that off-diagonal perturbation of a block diagonal matrix (or matrix polynomial) result in perturbation in the eigenvalues that scale quadratically with the norm of the perturbation. The quadratic scaling hinges on the block structure of the matrices as in (3) or (27), which the eigenvectors inherit. In fact, even tighter bounds can be obtained if further block structure is present, such as block tridiagonal [11]. In addition to some indicated in the text, possible future directions include investigating the accuracy in the expansion and residual bounds and in such cases, examine the implications in terms of the eigenvectors, and overcoming the case where the gap is too small for the bounds to be of use. Eigenvalue perturbation theory is a well-established yet active and useful area of research.

References

1. Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst. *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. SIAM, Philadelphia, PA, USA, 2000.
2. I. Gohberg, P. Lancaster, and L. Rodman. *Matrix polynomials*. SIAM, Philadelphia, USA, 2009. Unabridged republication of book first published by Academic Press in 1982.
3. G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, 4th edition, 2012.
4. N. J. Higham, R.-C. Li, and F. Tisseur. Backward error of polynomial eigenproblems solved by linearization. *SIAM J. Matrix Anal. Appl.*, 29(4):1218–1241, 2007.
5. W. Kahan, B. N. Parlett, and E. Jiang. Residual bounds on approximate eigensystems of nonnormal matrices. *SIAM J. Numer. Anal.*, 19(3):470–484, 1982.
6. C.-K. Li and R.-C. Li. A note on eigenvalues of perturbed Hermitian matrices. *Linear Algebra Appl.*, 395:183–190, 2005.
7. R.-C. Li. Matrix perturbation theory. In L. Hogben, R. Brualdi, A. Greenbaum, and R. Mathias, editors, *Handbook of Linear Algebra*, chapter 15. CRC Press, Boca Raton, FL, 2006.
8. R.-C. Li, Y. Nakatsukasa, N. Truhar, and S. Xu. Perturbation of partitioned Hermitian generalized eigenvalue problem. *SIAM J. Matrix Anal. Appl.*, 32(2):642–663, 2011.

9. R. Mathias. Quadratic residual bounds for the Hermitian eigenvalue problem. *SIAM J. Matrix Anal. Appl.*, 19(2):541–550, 1998.
10. Y. Nakatsukasa. Gerschgorin's theorem for generalized eigenvalue problems in the Euclidean metric. *Math. Comp.*, 80(276):2127–2142, 2011.
11. Y. Nakatsukasa. Eigenvalue perturbation bounds for Hermitian block tridiagonal matrices. *Appl. Numer. Math.*, 62(1):67 – 78, 2012.
12. B. N. Parlett. *The Symmetric Eigenvalue Problem*. SIAM, Philadelphia, 1998.
13. G. W. Stewart. *Matrix Algorithms Volume II: Eigensystems*. SIAM, Philadelphia, 2001.
14. G. W. Stewart and J.-G. Sun. *Matrix Perturbation Theory (Computer Science and Scientific Computing)*. Academic Press, 1990.
15. F. Tisseur. Backward error and condition of polynomial eigenvalue problems. *Linear Algebra Appl.*, 309(1):339–361, 2000.
16. R. S. Varga. *Geršgorin and His Circles*. Springer-Verlag, 2004.