COMPUTING ALL PAIRS (λ, μ) SUCH THAT λ IS A DOUBLE EIGENVALUE OF $A + \mu B$

ELIAS JARLEBRING*, SIMEN KVAAL†, AND WIM MICHIELS*

Abstract. Double eigenvalues are not generic for matrices without any particular structure. A matrix depending linearly on a scalar parameter, $A + \mu B$, will however generically have double eigenvalues for some values of the parameter μ. In this paper we consider the problem of finding those values. More precisely, we construct a method to accurately find *all* scalar pairs (λ, μ) such that $A + \mu B$ has a double eigenvalue λ, where A and B are given arbitrary complex matrices. The general idea of the globally convergent method is that if μ is close to a solution then $A + \mu B$ has two eigenvalues which are close to each other. We fix the relative distance between these two eigenvalues and construct a method to solve and study it by observing that the resulting problem can be stated as a *two-parameter eigenvalue problem*, which is already studied in the literature. The method, which we call the *method of fixed relative distance* (MFRD), involves solving a two-parameter eigenvalue problem which returns approximations of all solutions. It is unfortunately not possible to get full accuracy with MFRD. In order to compute solutions with full accuracy, we present an iterative method which returns a very accurate solution, for a sufficiently good starting value. The approach is illustrated with one academic example and one application to a simple problem in computational quantum mechanics.

1. Introduction. For two arbitrary, given matrices $A, B \in \mathbb{C}^{n \times n}$, we will look for pairs $(\lambda, \mu) \in \mathbb{C}^2$ such that the matrix $A + \mu B$ has a *double* eigenvalue at λ. More precisely, we will consider this problem and present a globally convergent, accurate method for all the solution pairs (λ, μ). The set of pairs (λ, μ) for which λ is an eigenvalue of $A + \mu B$ of multiplicity two or greater will be denoted by $\mathbb{B}(A, B)$, i.e.,

$$\mathbb{B}(A, B) := \{(\lambda, \mu) \in \mathbb{C}^2 : \lambda \text{ is a non-simple eigenvalue of } A + \mu B\}.$$

In the same way that double eigenvalues for matrices without a particular structure are not generic, we also have that an element of $\mathbb{B}(A, B)$ corresponding to an eigenvalue of multiplicity greater than two is not generic. That is, λ is generically an eigenvalue of $A + \mu B$ with multiplicity two if $(\lambda, \mu) \in \mathbb{B}(A, B)$. We will see (in Section 2) that the degenerate situation corresponds to a badly conditioned problem and there is little hope to construct a very accurate algorithm. We will for this reason focus the study on the generic situation where λ is a double eigenvalue of $A + \mu B$, although the algorithm is applicable to the degenerate case.

Our study is in a sense a simultaneous numerical characterization of the perturbation of several double eigenvalues for a structured perturbation. Repeated eigenvalues have received a lot of attention in the field of perturbation theory. See, e.g., the standard works [16, 5, 13]. Despite the fact that the research field related to repeated eigenvalues is quite mature, there appears to be no globally convergent numerical method for the problem of finding all elements of $\mathbb{B}(A, B)$.

There exist, for instance, the results on generalizations of condition numbers for repeated eigenvalues [17, 22, 28, 15, 7, 6], which provide a theory for small perturbations of the repeated eigenvalue. Note that these results do not provide information about all the elements of $\mathbb{B}(A, B)$. In our problem formulation, μ is generally not small, and the set $\mathbb{B}(A, B)$ is not well approximated by an asymptotic characterization for μ close to zero.

---

*K.U. Leuven, Department of Computer Science, 3001 Heverlee, Belgium, {`Elias.Jarlebring,` `Wim.Michiels`}`@cs.kuleuven.be`

†CMA, P.O. Box 1053 Blindern, NO-0316 Oslo, Norway, `simen.kvaal@cma.uio.no`

There is also recent research on the *Wilkinson distance* [2, 26, 9, 18, 19, 20], which is defined as the size of the smallest perturbation E such that $A + E$ has a multiple eigenvalue. Suppose the direction of the perturbation is known, i.e., $E = \mu B$, where B is a known matrix and μ is unknown. It is easy to compute this structured Wilkinson distance with a method which computes $\mathbb{B}(A, B)$ by choosing the smallest μ in $\mathbb{B}(A, B)$. Note, however, that with the methods in this paper we can compute all elements of $\mathbb{B}(A, B)$ and not only the smallest μ. Moreover, although there are estimates and numerical methods for the Wilkinson distance (see e.g. the references in [2]), they do not seem to be easily restricted or extended to this structured case. In particular, there exist local iterative methods such as [18]. The method in [18] is for an unstructured problem and designed for repeated eigenvalues with an a priori known Jordan structure. There is also a method in [19] which is a min-max formula for the Wilkinson distance. Note that the optimization problem is possibly difficult to solve in a reliable way. We present a method which is globally convergent and only involves linear algebra operations.

For general matrices, double eigenvalues correspond to a degenerate case and the study of double eigenvalues has traditionally been somewhat theoretically oriented. However, the results of this paper were initiated and motivated by an application. In quantum mechanical perturbation theory, the radius of convergence of a perturbation series in powers of μ for eigenvalues of $A + \mu B$ are determined by certain points in $\mathbb{B}(A, B)$ [27]. These points are of course not known for realistic problems, and a method to actually estimate the radius of convergence for different perturbations is of great utility. The points where there is a repeated eigenvalue are also used in the analysis of root loci in control theory. For other applications of repeated eigenvalues see the references in [18].

In order to outline the contents of the paper we first briefly introduce some terminology. Perturbations of double eigenvalues are considerably different from the theory of simple eigenvalues. For instance, the eigenvalues typically do not behave in an analytic way. Throughout this paper we will make use of the fact that λ as a function of μ is an algebraic function and can be expanded in a Puiseux series [16, Chapter II, Section 1.2.]

$$\lambda(\mu + \Delta) = \lambda(\mu) + c\Delta^\alpha + o(\Delta^\alpha), \tag{1.1}$$

where α is a rational number and μ such that $\lambda(\mu)$ is a repeated eigenvalue. For double eigenvalues we have $\alpha \in \frac{1}{2}\mathbb{Z}$ and c can be chosen non-zero. If α is not an integer, then (1.1) characterizes a branch point of the eigenvalue map $\lambda(\mu)$.

Before presenting the numerical schemes, we prove some properties of the problem (in Section 2). It turns out that the problem is numerically well posed in the generic case. We motivate this by proving preservation of existence of solutions for infinitesimal perturbations as well as continuity and conditioning properties. It turns out that if λ is an eigenvalue of $A + \mu B$ with multiplicity higher than two, the problem is not well conditioned. Although the theory only supports the generic case that λ is a double eigenvalue, the construction of the methods is, however, such that they are expected to work also in these situations.

It follows from the continuity of eigenvalue paths that if a matrix is a sufficiently small perturbation of a matrix having a double eigenvalue, it will have two eigenvalues close to each other. In the method presented in Section 3, we will use this property to construct a globally convergent method, based on fixing the relative distance between two eigenvalues to a small value. This turns out to be equivalent to a so-called *two-*

parameter eigenvalue problem and can be solved with methods in the literature. We call it the *method of fixed relative distance* (MFRD).

The construction of the method in Section 3 (MFRD) is such that it will not yield an exact solution; not even in exact arithmetic. Hence, in practice, the solution will not be of full precision. Since we wish to have an accurate solution, an iterative Newton type method for elements of $\mathbb{B}(A, B)$ is presented (in Section 4). The starting value of the iterative method is taken as the result of MFRD. The combined algorithm is globally convergent and produces accurate solutions. We illustrate the methods and theory with some examples in Section 5.

2. Properties of the problem. In later sections we will construct numerical algorithms for $\mathbb{B}(A, B)$. One can only expect numerically relevant results if the problem is posed in a numerically reasonable way. For this reason, we will in this section demonstrate some important properties of $\mathbb{B}(A, B)$. Without these, one could not expect to reliably solve the problem numerically. We study
  • the cardinality of $\mathbb{B}(A, B)$;
  • the continuous preservation of elements under perturbation; and
  • the conditioning of the problem.
The conclusion of the study is that the elements of $\mathbb{B}(A, B)$ are generically finite, that the elements are always preserved under sufficiently small perturbations and that the sensitivity is generically finite.

In the technical reasoning we will make use of the property that if $(\lambda, \mu) \in \mathbb{B}(A, B)$ then μ is explicitly given as the root of a polynomial. To this end let

$$f(\lambda, \mu; \ A, B) := \det(\lambda I - A - \mu B). \tag{2.1}$$

Since λ is a double eigenvalue, a pair $(\lambda, \mu) \in \mathbb{B}(A, B)$ can now equivalently be characterized as

$$f(\lambda, \mu; \ A, B) = 0 \tag{2.2a}$$
$$f_\lambda(\lambda, \mu; \ A, B) = 0, \tag{2.2b}$$

where f_λ is the short notation for $\frac{\partial f}{\partial \lambda}(\lambda, \mu; \ A, B)$. Let $\mu \mapsto M_1(\mu; \ A, B) \in \mathbb{C}^{n \times n}$ respectively $\mu \mapsto M_2(\mu; \ A, B) \in \mathbb{C}^{(n-1) \times (n-1)}$ represent the companion matrices corresponding to the polynomial $f(\cdot, \mu; \ A, B)$, respectively $f_\lambda(\cdot, \mu; \ A, B)$. Then we can write (2.2) as

$$\lambda \in \sigma(M_1(\mu; \ A, B)) \tag{2.3a}$$
$$-\lambda \in \sigma(-M_2(\mu; \ A, B)), \tag{2.3b}$$

which allows us to eliminate one of the variables, say λ.

PROPOSITION 2.1 (Explicit form). *Let $\mu_* \in \mathbb{C}$. The following assertions are equivalent.*
  1. $(\lambda, \mu_*) \in \mathbb{B}(A, B)$ for some $\lambda \in \mathbb{C}$.
  2. μ_* is a solution of the polynomial eigenvalue problem

$$\det\left(M_1\left(\mu; \ A, B\right) \oplus \left(-M_2\left(\mu; \ A, B\right)\right)\right) = 0. \tag{2.4}$$

Here, we have as usual (in e.g. [13]) denoted the Kronecker sum by $\oplus$. The Kronecker sum is defined as $A \oplus B = A \otimes I_1 + I_2 \otimes B$, where I_1, I_2 are the identity matrices of appropriate size. We have used the property of the Kronecker sum that any eigenvalue of $A \oplus B$ is the sum of an eigenvalue of A and an eigenvalue of B.

REMARK 2.2 (Numerical stability of the explicit form). *Proposition 2.1 directly gives rise to a conceptual algorithm for the computation of $\mathbb{B}(A, B)$, which consists of first solving (2.4) for μ, and, next, computing the corresponding values of λ from (2.3). From a numerical point of view this approach is to be avoided as it requires the explicit computation of the scalar characteristic equation, i.e., the symbolic computation of the determinant (in (2.1)) and the derivative.*

2.1. Cardinality. It is not reasonable to expect that a numerical algorithm can find all solutions of a problem if the problem has an infinite number of solutions. Fortunately, the generic situation is that $\mathbb{B}(A, B)$ only contains a finite number of elements.

PROPOSITION 2.3 (Cardinality). *The following classification holds.*

   1. If the nondegeneracy condition

$$\det(M_1(\mu; \ A, B) \oplus (-M_2(\mu; \ A, B))) \not\equiv 0$$

   is satisfied, then the set $\mathbb{B}(A, B)$ consists of a finite number of isolated pairs (λ, μ) in $\mathbb{C} \times \mathbb{C}$.
   2. Conversely, if

$$\det(M_1(\mu; \ A, B) \oplus (-M_2(\mu; \ A, B))) \equiv 0, \qquad (2.5)$$

   then for all $\mu \in \mathbb{C}$ there exists a $\lambda \in \mathbb{C}$ such that $(\lambda, \mu) \in \mathbb{B}(A, B)$.

In this work we will focus on the first case since it is generic and the second case can be easily handled in practice by computing the eigenvalues of $A + \mu B$ for several μ.

2.2. Continuity. A numerical algorithm will always introduce rounding errors. If the output is not continuous with respect to the input or (more critically) solutions appear or disappear under sufficiently small perturbations, there is little hope to construct a robust numerical scheme. The problem we are considering fulfills this necessary condition, as the set $\mathbb{B}(A, B)$ is continuous with respect to changes in A and B in the following sense.

PROPOSITION 2.4 (Continuity). *Let (λ_*, μ_*) be an isolated pair of $\mathbb{B}(A, B)$. Then there exists a number $\hat{\gamma} > 0$ such that for all $\gamma \in (0, \hat{\gamma})$ there is a number $\delta > 0$ such that $\mathbb{B}(A + \Delta A, B + \Delta B)$ contains at least one pair (λ, μ) satisfying $|\lambda - \lambda_*| < \gamma$ and $|\mu - \mu_*| < \gamma$, whenever $\|\Delta A\| < \delta$ and $\|\Delta B\| < \delta$.*

Proof. The proof follows from the continuous dependence of the solutions of (2.4), a one-parameter polynomial eigenvalue problem, with respect to A and B, combined with the continuous dependence of the zeros of $f(\cdot, \mu; \ A, B)$ with respect to A, B and μ. $\square$

2.3. Conditioning. Continuity with respect to changes in the input is not sufficient for the problem to be numerically well posed. If the output is highly sensitive to perturbations in the input, the problem is also considered very difficult from a numerical point of view. This comes from the fact that the first operation in a numerical algorithm will introduce rounding errors. In the following we present a characterization of the case that such ill-conditioning, i.e., high sensitivity with respect to input, occurs.

Let (λ_*, μ_*) be an isolated pair of $\mathbb{B}(A, B)$ and consider the corresponding solutions of $\mathbb{B}(A + \epsilon E_a, B + \epsilon E_b)$, where $\epsilon > 0$ is a small perturbation parameter and E_a

and E_b are n-by-n arbitrary complex matrices. The characterization (2.2) brings us to the equations

$$f(\lambda_* + \Delta\lambda, \mu_* + \Delta\mu; \ A + \epsilon E_a, B + \epsilon E_b) = 0 \qquad (2.6a)$$

$$f_\lambda(\lambda_* + \Delta\lambda, \mu_* + \Delta\mu; \ A + \epsilon E_a, B + \epsilon E_b) = 0, \qquad (2.6b)$$

where we made the substitutions $\lambda = \lambda_* + \Delta\lambda$ and $\mu = \mu_* + \Delta\mu$. From the implicit function theorem, we conclude that if the Jacobian matrix

$$C(\lambda_*, \mu_*) := \begin{bmatrix} 0 & f_\mu(\lambda_*, \mu_*; \ A, B) \\ f_{\lambda\lambda}(\lambda_*, \mu_*; \ A, B) & f_{\lambda\mu}(\lambda_*, \mu_*; \ A, B) \end{bmatrix} \qquad (2.7)$$

is invertible, then (2.6) locally defines a unique function $\epsilon \mapsto (\Delta\lambda(\epsilon), \Delta\mu(\epsilon))$ that can be expanded as

$$\begin{bmatrix} \Delta\lambda \\ \Delta\mu \end{bmatrix} = -C(\lambda_*, \mu_*)^{-1} \begin{bmatrix} \frac{\partial}{\partial\epsilon} f(\lambda_*, \mu_*; \ A + \epsilon E_a, B + \epsilon E_b)\big|_{\epsilon=0} \\ \frac{\partial}{\partial\epsilon} f_\lambda(\lambda_*, \mu_*; \ A + \epsilon E_a, B + \epsilon E_b)\big|_{\epsilon=0} \end{bmatrix} \epsilon + \mathcal{O}(\epsilon^2). \quad (2.8)$$

We note that the matrix $C(\lambda_*, \mu_*)$ is invertible if and only if

$$f_{\lambda\lambda}(\lambda_*, \mu_*) f_\mu(\lambda_*, \mu_*) \neq 0. \qquad (2.9)$$

This condition corresponds to the generic situation where λ_* is a double, non-semisimple eigenvalue of $A + \mu_* B$ that satisfies the *completely regular splitting property*. The local behavior of a perturbation is said to have a completely regular splitting if the order of the root in the first non-vanishing term in the Puiseux series of the eigenvalue coincides with the partial multiplicities in the Jordan structure. See, e.g., [14] and the references therein for literature on completely regular splitting. Here we will only use the fact that complete regular splitting is the generic case. We summarize it as follows.

PROPOSITION 2.5. *Let (λ_*, μ_*) be an isolated pair of $\mathbb{B}(A, B)$.*
  1. *If (2.9) holds, then the sensitivity of the pair (λ_*, μ_*) is determined by*

$$\|C(\lambda_*, \mu_*)^{-1}\|,$$

  where $C(\lambda_*, \mu_*)$ is defined by (2.7). Furthermore, the expansion (2.8) is applicable.
  2. *The pair (λ_*, μ_*) is ill conditioned if $f_{\lambda\lambda}(\lambda_*, \mu_*) = 0$, i.e., the eigenvalue λ_* of $A + \mu_* B$ has multiplicity larger than two.*
  3. *The pair (λ_*, μ_*) is ill conditioned if $f_\mu(\lambda_*, \mu_*) = 0$. This includes the case where λ_* is a double semisimple eigenvalue of $A + \mu_* B$.*

3. A method of fixed relative distance. Consider for the moment a fixed scalar $\mu \in \mathbb{C}$. Suppose $\lambda \in \mathbb{C}$ is a complex scalar fulfilling two conditions; it is an eigenvalue of $A + \mu B$, i.e.,

$$\lambda \in \sigma(A + \mu B) \qquad (3.1)$$

and $(1 + \varepsilon)\lambda$ is also an eigenvalue of $A + \mu B$, i.e.,

$$(1 + \varepsilon)\lambda \in \sigma(A + \mu B), \qquad (3.2)$$

for a fixed non-zero complex scalar ε. The fundamental idea of the method we will present in this section is that the solutions of (3.1) and (3.2), where λ and μ are the unknowns, approximate the elements of $\mathbb{B}(A,B)$ for sufficiently small ε. We will denote the set of all solutions of (3.1) and (3.2) by $\mathbb{D}(A,B,\varepsilon)$, i.e.,

$$\mathbb{D}(A,B,\varepsilon) := \{\ (\lambda,\mu) \in \mathbb{C}^2 : (3.1) \text{ and } (3.2) \text{ hold } \}.$$

Note that if $(\lambda,\mu) \in \mathbb{D}(A,B,\varepsilon)$, then the matrix $A+\mu B$ has two eigenvalues with relative distance ε. The conceptual method to estimate $\mathbb{B}(A,B)$ by computing $\mathbb{D}(A,B,\varepsilon)$ for sufficiently small ε will be referred to as the *method of fixed relative distance* (MFRD).

We will first see that $\mathbb{D}(A,B,\varepsilon)$ can be computed by solving a generalized eigenvalue problem and then we will study the elements of $\mathbb{D}(A,B,\varepsilon)$ as a function of ε. In Section 3.1 we show that the finite limits (as $\varepsilon \to 0$) correspond to the elements of $\mathbb{B}(A,B)$. We study the case where an element is unbounded as $\varepsilon \to 0$ in Section 3.2. The error of the approximation is characterized in Section 3.3 and used to find a reasonable choice of ε in Section 3.4.

In Proposition 2.1 we saw that if $(\lambda,\mu) \in \mathbb{B}(A,B)$, then μ was a root of a polynomial. The direct method to compute the roots of this polynomial is not attractive from a the point of view of numerical stability. This is resolved with the regularized formulation of $\mathbb{D}(A,B,\varepsilon)$. We will now see that determining $\mathbb{D}(A,B,\varepsilon)$ is a problem of the type called *two-parameter eigenvalue problems*, which are solvable in a numerically stable way. The relations (3.1) and (3.2) can be reformulated as follows.

PROBLEM 3.1 (The associated two-parameter eigenvalue problem). *Given $A, B \in \mathbb{C}^{n\times n}$, $\varepsilon \in \mathbb{C}$, find $(\lambda,\mu) \in \mathbb{C}^2$ and a pair of non-trivial vectors $(u,v) \in \mathbb{C}^{n\times 2}$ such that*

$$Au = (\lambda I - \mu B)u,$$
$$Av = (\lambda(1+\varepsilon)I - \mu B)v.$$

There is a general theory for the two-parameter eigenvalue problem available in the classical works [4, 3]. There are also more recent results which are more numerically oriented [12]. Some results are focused on the singular problem [24] and some are for different types of generalizations [23]. We also note that multiparameter eigenvalue problems have been used for (relative) placement of eigenvalues [8].

There are several ways to numerically solve two-parameter eigenvalue problems, e.g., the Jacobi-Davidson type method in [11]. For further methods, see the references in [11]. The most common way to solve and analyze two-parameter eigenvalue problems is by means of three matrices $\Delta_0, \Delta_1, \Delta_2 \in \mathbb{C}^{n^2 \times n^2}$, called the matrix determinants. In the case of Problem 3.1, the matrix determinants are,

$$\Delta_0(\varepsilon) = -I \otimes B + (1+\varepsilon)B \otimes I = ((1+\varepsilon)B) \oplus (-B), \tag{3.3}$$
$$\Delta_1 = -A \otimes B + B \otimes A, \tag{3.4}$$
$$\Delta_2(\varepsilon) = I \otimes A - (1+\varepsilon)A \otimes I = (-(1+\varepsilon)A) \oplus A. \tag{3.5}$$

One reason why the matrix determinants are important in the context of two-parameter eigenvalue problems, stems from the fact that the two-parameter eigenvalue problem is (under sufficient non-singularity conditions) equivalent to the solutions of two coupled generalized eigenvalue problems. In this case

$$\lambda(\varepsilon)\Delta_0(\varepsilon)z = \Delta_1 z \tag{3.6}$$

and

$$\mu(\varepsilon)\Delta_0(\varepsilon)z = \Delta_2(\varepsilon)z, \tag{3.7}$$

where z is a decomposable tensor $z = u \otimes v$. Although not always very efficient, we will solve the relative distance problem by solving (3.6) and (3.7). It follows from the generalized eigenvalue problems (3.6) and (3.7) that the set $\mathbb{D}(A, B, \varepsilon)$ is (generically) a union of a finite number of functions in ε. For this reason we have denoted the solutions by $\lambda(\varepsilon)$ and $\mu(\varepsilon)$ when we wish to stress the dependence on ε. Note that since all eigenvalue paths are continuous (but not necessarily differentiable), the elements of $\mathbb{D}(A, B, \varepsilon)$ are also continuous with respect to ε.

A two-parameter eigenvalue problem for which $\Delta_0(\varepsilon)$ is non-singular is completely characterized by the eigenvalues of $\Delta_0(\varepsilon)^{-1}\Delta_1$ and $\Delta_0(\varepsilon)^{-1}\Delta_2(\varepsilon)$, and hence somewhat easier from a theoretical and numerical perspective. The condition that $\Delta_0(\varepsilon)$ is singular is one way to define singularity of the two-parameter eigenvalue (also used e.g. in [24]). The two-parameter eigenvalue problem in Problem 3.1 is in this sense indeed generically non-singular for sufficiently small non-zero $\varepsilon \in \mathbb{C}\backslash\{0\}$, as can be seen in the following lemma.

LEMMA 3.2 (Non-singularity). *The two-parameter eigenvalue problem (Problem 3.1) is singular, i.e.,* $\det(\Delta_0(\varepsilon)) = 0$, *iff there is* $\gamma \in \sigma(B)$ *such that* $\gamma(1 + \varepsilon) \in \sigma(B)$.

REMARK 3.3 (The non-zero regularization parameter ε). *If ε is set to zero, it is easy to see from Lemma 3.2 that the two-parameter eigenvalue problem is singular. Although there are results on singular two-parameter eigenvalue problems [24] such an approach seems impossible for this problem. Note that if $\varepsilon = 0$, then the two eigenvalue problems (3.1) and (3.2) are identical and the solution corresponds to the eigenvalue curves $\lambda(\mu) \in \sigma(A + \mu B)$. The points on these curves do, generically, not correspond to double eigenvalues.*

The two-parameter eigenvalue problem for $\varepsilon = 0$ is singular and the corresponding problem for $\mathbb{D}(A, B, \varepsilon)$ is not singular. Hence, the problem associated with MFRD is in a sense a regularized problem and we will call the parameter ε a regularization parameter.

In the remaining parts of this section, we wish to show properties of $\mathbb{D}(A, B, \varepsilon)$ and in particular in what sense $\mathbb{D}(A, B, \varepsilon)$ approximates $\mathbb{B}(A, B)$.

3.1. Consistency. The first property which we will illustrate in a remark and a theorem is that the method is consistent, in the sense that (generically) all finite limits of $\mathbb{D}(A, B, \varepsilon)$ as $\varepsilon \to 0$ belong to $\mathbb{B}(A, B)$. Moreover, this limit is independent of the angle of ε.

REMARK 3.4 (The angle of ε). *Since ε is the relative distance, fixing the complex angle of ε fixes the angle of the separation of the eigenvalues $\lambda(\varepsilon)$ and $(1 + \varepsilon)\lambda(\varepsilon)$. We will now see that the angle of ε will asymptotically only influence the angle of the approximation error. The asymptotic effect of changing φ where $\varepsilon = |\varepsilon|e^{i\varphi}$ can be roughly motivated as follows. Consider the first terms in a completely regular square root splitting, $\lambda(\mu(\varepsilon)) = \lambda_* \pm \sqrt{\mu_* - \mu(\varepsilon)}$. The condition that the relative distance is ε implies that $\mu_* - \mu(\varepsilon) = c\varepsilon^2 = c|\varepsilon|^2e^{i2\varphi}$. This in turn implies that the error of the approximation is given by $\lambda(\mu(\varepsilon)) - \lambda_* = \pm\sqrt{\mu_* - \mu(\varepsilon)} = \pm c^{1/2}|\varepsilon|e^{\pm i\varphi}$. Now note that φ, i.e., the angle of the relative distance ε, only changes the angle of the error and not the magnitude. The phenomenon is illustrated in Figure 3.1. The same reasoning holds for an arbitrary splitting of the eigenvalue.*
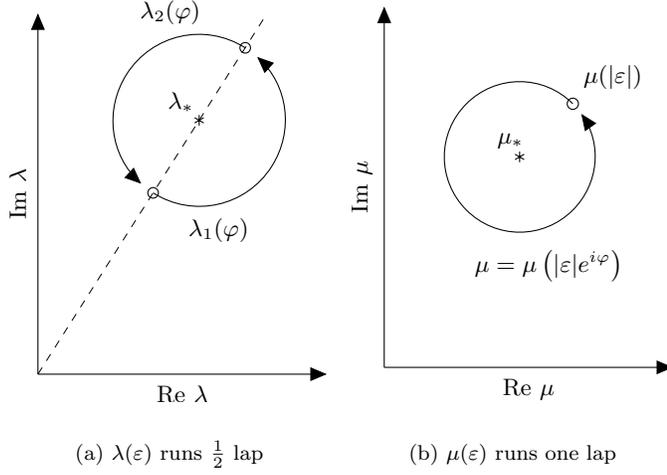
(a) $\lambda(\varepsilon)$ runs $\frac{1}{2}$ lap   (b) $\mu(\varepsilon)$ runs one lap

FIGURE 3.1. *The approximations rotate around the solutions λ_* and μ_* as the angle of ε is changed. The plots show the curves for $\varphi \in [0, \pi]$. We have denoted $\lambda_1(\varphi) := \lambda\left(|\varepsilon|e^{i\varphi}\right)$ and $\lambda_2(\varphi) := \lambda_1(\varphi)(1 + |\varepsilon|e^{i\varphi})$. The dashed line corresponds to $\varphi = 0$. In subfigure (a), the same dashed line (asymptotically) corresponds to the line to the origin.*

In the following theorem we formalize the argument above and see that the set of finite limits of $\mathbb{D}(A, B, \varepsilon)$ as $\varepsilon \to 0$ is generically equal to the set $\mathbb{B}(A, B)$.

THEOREM 3.5 (Consistency). *Consider the pair $A, B \in \mathbb{C}^{n \times n}$. The sets $\mathbb{D}(A, B, \varepsilon)$ and $\mathbb{B}(A, B)$ are related by the following statements:*

    i) If B is non-singular, the set $\mathbb{D}(A, B, \varepsilon)$ is the union of n^2 pair of functions $(\lambda(\varepsilon), \mu(\varepsilon))$ for sufficiently small $\varepsilon \neq 0$.

    ii) If $(\lambda(\varepsilon), \mu(\varepsilon)) \in \mathbb{D}(A, B, \varepsilon)$ and $(\lambda(\varepsilon), \mu(\varepsilon)) \to (\lambda_*, \mu_*) \in \mathbb{C}^2$ as $|\varepsilon| \to 0$ and $\lambda_* \neq 0$ then $(\lambda_*, \mu_*) \in \mathbb{B}(A, B)$.

    iii) If $(\lambda_*, \mu_*) \in \mathbb{B}(A, B)$ is an isolated pair and $\lambda_* \neq 0$ then there exists a path $\varepsilon \mapsto (\lambda(\varepsilon), \mu(\varepsilon))$ such that $(\lambda(\varepsilon), \mu(\varepsilon)) \in \mathbb{D}(A, B, \varepsilon)$ for every $\varepsilon > 0$ in a neighborhood of $\varepsilon = 0$ and such that $(\lambda(\varepsilon), \mu(\varepsilon)) \to (\lambda_*, \mu_*)$ as $\varepsilon \to 0$.

Proof. The paths $\lambda(\varepsilon)$ and $\mu(\varepsilon)$ are solutions of the generalized eigenvalue problems corresponding to the pencils $\Delta_1 - \lambda(\varepsilon)\Delta_0(\varepsilon)$ and $\Delta_2(\varepsilon) - \mu(\varepsilon)\Delta_0(\varepsilon)$. The eigenvalues of a matrix depending continuously on a parameter is a collection of paths, where the number of paths equals the dimension of the matrix, here n^2 (see e.g. [10, Corollary 4.2.4]). The generalized eigenvalue problems can be rewritten as standard eigenvalue problems if $\Delta_0(\varepsilon)$ is non-singular. Hence, to show i) it is sufficient to show that $\Delta_0(\varepsilon)$ is non-singular for sufficiently small ε.

It follows from the sum-property of Kronecker sums that the eigenvalues of $\Delta_0(\varepsilon)$ are $(1 + \varepsilon)b_i - b_j$ for $i, j = 1, \ldots n$ where b_i are the eigenvalues of B. Since $b_i \neq 0$, all eigenvalues of $\Delta_0(\varepsilon)$ are non-zero for sufficiently small $\varepsilon \neq 0$. This proves i).

Let $f(\lambda, \mu) = \det(-\lambda I + A + \mu B)$. In order to show ii) we will show that $f_\lambda(\lambda_*, \mu_*) = 0$. First note that

$$f(\lambda(\varepsilon), \mu(\varepsilon)) = f(\lambda(\varepsilon)(1 + \varepsilon), \mu(\varepsilon)) = 0.$$

Hence, by Taylor expansion,

$$0 = f(\lambda(\varepsilon)(1 + \varepsilon), \mu(\varepsilon)) = f(\lambda(\varepsilon), \mu(\varepsilon)) + \lambda(\varepsilon)\varepsilon f_\lambda(\lambda(\varepsilon), \mu(\varepsilon)) + \mathcal{O}(\lambda(\varepsilon)\varepsilon)^2.$$

Since $\lambda(\varepsilon) \to \lambda_* \in \mathbb{C}\backslash\{0\}$ when $\varepsilon \to 0$, we have that $f_\lambda(\lambda_*, \mu_*) = 0$.

In order to show part iii), let $\lambda_0(r, \theta)$ and $\lambda_1(r, \theta)$ be two eigenvalues of $A + (\mu_* + re^{i\theta})B$ with $\lambda_i(0, 0) = \lambda_*$. We have an expansion

$$\lambda_1(r, \theta) = \lambda_* + Cr^\alpha e^{i\theta\alpha} + o(r^\alpha), \tag{3.8}$$

where $C \neq 0$ and $\alpha > 0$, and a similar expansion for $\lambda_0(r, \theta)$. Since the double eigenvalue is isolated, the difference does not vanish identically, and we have an expansion

$$\lambda_1(r, \theta) - \lambda_0(r, \theta) = Dr^\beta e^{i\theta\beta} + o(r^\beta), \tag{3.9}$$

for some $D \neq 0$. (Note that λ_i can be chosen as different branches from one expansion or from different Puiseux expansions.)

We first construct a function $\theta(r)$ such that the relative distance function $g(r, \theta(r))$ is real,

$$g(r, \theta(r)) := \frac{\lambda_1(r, \theta(r)) - \lambda_0(r, \theta(r))}{\lambda_0(r, \theta(r))} \in \mathbb{R}.$$

The angle of the relative distance is

$$\arg(g(r, \theta)) = \arg(D) + \theta\beta + o(r^\beta) - \arg(\lambda_*) + o(r^\alpha). \tag{3.10}$$

If $\arg(g(r, \theta)) = 0$, then $g(r, \theta) \in \mathbb{R}$. We use the implicit function theorem on this condition and $\frac{\partial}{\partial\theta}(\arg(g(r, \theta)))_{r=0} = \beta \neq 0$, which shows that there is a function $\theta(r)$, such that $g(r, \theta(r))$ is real for every r in a neighborhood of $r = 0$. Without loss of generality, we can impose that $g(r, \theta(r))$ is positive since switching λ_1 and λ_0 corresponds to switching sign of $g(r, \theta(r))$ for $r \to 0$.

We now use $1/(1 + x) = 1 - x + o(x)$, where $x = \lambda_*^{-1}\lambda_0(r, \theta(r)) - 1$ and the expansion (3.9) and find that,

$$g(r, \theta(r)) = \lambda_*^{-1}Dr^\beta e^{i\theta(r)\beta} + o(r^\beta) = |\lambda_*|^{-1}|D|r^\beta + o(r^\beta) \tag{3.11}$$

where in the last step we used the explicit formula for $\theta(r)$ from (3.10). Now consider,

$$h(r, \varepsilon) := \varepsilon^{1/\beta} - \left(\frac{\lambda_1(r, \theta(r)) - \lambda_0(r, \theta(r))}{\lambda_0(r, \theta(r))}\right)^{1/\beta} = \varepsilon^{1/\beta} - |\lambda_*|^{-1/\beta}|D|^{1/\beta}r + o(r). \tag{3.12}$$

We again use the implicit function theorem; now on $h(r, \varepsilon)$. Since $\frac{\partial h}{\partial r}(0, 0) \neq 0$, there exists a function $r(\varepsilon)$ such that

$$\varepsilon = \frac{\lambda_1(r(\varepsilon), \theta(r(\varepsilon))) - \lambda_0(r(\varepsilon), \theta(r(\varepsilon)))}{\lambda_0(r, \theta(r(\varepsilon)))}, \tag{3.13}$$

for every $\varepsilon > 0$ in a neighborhood of $\varepsilon = 0$. By construction, $\lambda_0(r(\varepsilon), \theta(r(\varepsilon)))$ and $\lambda_1(r(\varepsilon), \theta(r(\varepsilon)))$ are eigenvalues of $A + (\mu_* + r(\varepsilon)e^{i\theta(r(\varepsilon))})B$ and from (3.13) the relative distance is ε. Hence, $\varepsilon \mapsto (\lambda_1(r(\varepsilon), \theta(r(\varepsilon))), \mu_* + r(\varepsilon)e^{i\theta(r(\varepsilon))})$ is a function corresponding to an element of $\mathbb{D}(A, B, \varepsilon)$. $\square$

REMARK 3.6 ($\lambda_* = 0$). *Note that the case $\lambda_* = 0$ has to be explicitly excluded in Theorem 3.5. This stems from the fact that the relative distance is defined by (3.1) and (3.2), which are trivially fulfilled if $\lambda = 0$ (and if μ is an eigenvalue of the generalized eigenvalue problem $A + \mu B$) for any ε. The special case where $\lambda = 0$ is a double eigenvalue can be easily handled by hand, and the solutions corresponding to $\lambda = 0$ can be safely and easily excluded from the solution set in the implementation.*

3.2. Unbounded limits. We now know (from Theorem 3.5) that the finite limits of $\mathbb{D}(A, B, \varepsilon)$ form, in a relevant way, the set $\mathbb{B}(A, B)$. At this point, it is important to note that this does not necessarily imply that the elements of $\mathbb{D}(A, B, \varepsilon)$ always approximate a corresponding element of $\mathbb{B}(A, B)$ as $\varepsilon \to 0$. It turns out that some elements of $\mathbb{D}(A, B, \varepsilon)$ can be unbounded as $\varepsilon \to 0$. This is illustrated in Theorem 3.8, for which we first need the following result which relates $\mathbb{D}(A, B, \varepsilon)$ with the converse problem $\mathbb{D}(B, A, \varepsilon)$.

LEMMA 3.7 (A converse identity). *Let $\varepsilon, \mu \in \mathbb{C} \setminus \{0\}$ be given. Then,*

$$(\lambda, \mu) \in \mathbb{D}(A, B, \varepsilon) \Leftrightarrow (\lambda/\mu, 1/\mu) \in \mathbb{D}(B, A, \varepsilon).$$

Proof. Define $f(\lambda, \mu) = \det(A + \mu B - \lambda I)$ and $g(\gamma, \nu) = \det(B + \nu A - \gamma I)$. Now, $(\lambda, \mu) \in \mathbb{D}(A, B, \varepsilon)$ if and only if $f(\lambda, \mu) = f((1+\varepsilon)\lambda, \mu) = 0$ if and only if $g(\lambda/\mu, 1/\mu) = g((1+\varepsilon)\lambda/\mu, 1/\mu) = 0$, since $f(\mu, \lambda) = \mu^n g(\lambda/\mu, 1/\mu)$. Since the relative distance between λ/μ and $(1+\varepsilon)\lambda/\mu$ is ε, $(\lambda/\mu, 1/\mu) \in \mathbb{D}(B, A, \varepsilon)$ as claimed. □

THEOREM 3.8 (Unbounded elements). *Suppose B is invertible and $\mathbb{B}(A, B)$ only consists of isolated elements. Let $\gamma_* \in \mathbb{C} \setminus \{0\}$. Then these statements are equivalent:*
  1. γ_* is a multiple eigenvalue of B
  2. There exists a punctured disc $\mathcal{D} = \{\varepsilon : |\varepsilon| < \hat{\varepsilon}\} \setminus \{0\} \subset \mathbb{C}$ and functions

$$\varepsilon \in \mathcal{D} \mapsto (\lambda(\varepsilon), \mu(\varepsilon)) \in \mathbb{C} \times \mathbb{C}$$

  satisfying $(\lambda(\varepsilon), \mu(\varepsilon)) \in \mathbb{D}(A, B, \varepsilon)$ for all $\varepsilon \in \mathcal{D}$ with

$$\lim_{\varepsilon \to 0} |\lambda(\varepsilon)| = \lim_{\varepsilon \to 0} |\mu(\varepsilon)| = \infty,$$

  and

$$\lim_{\varepsilon \to 0} \frac{\lambda(\varepsilon)}{\mu(\varepsilon)} = \gamma_*.$$

Proof. Let $f(\lambda, \mu)$ and $g(\gamma, \nu)$ be as in Lemma 3.7. We first prove that statement 1 implies statement 2. Since B is invertible, we have $\gamma_* \neq 0$. By point *(iii)* of Theorem 3.5, there exists a small punctured disk $\mathcal{D}$ and functions $\gamma, \nu : \mathcal{D} \to \mathbb{C} \times \mathbb{C}$ with $(\gamma(\varepsilon), \nu(\varepsilon)) \in \mathbb{D}(B, A, \varepsilon)$ and $\lim_{\varepsilon \to 0}(\gamma(\varepsilon), \nu(\varepsilon)) = (\gamma_*, 0)$.

By Lemma 3.7,

$$(\lambda(\varepsilon), \mu(\varepsilon)) := \left(\frac{\gamma(\varepsilon)}{\nu(\varepsilon)}, \frac{1}{\nu(\varepsilon)}\right) \in \mathbb{D}(A, B, \varepsilon) \quad \forall \varepsilon \in \mathcal{D}.$$

Moreover,

$$\lim_{\varepsilon \to 0} |\lambda(\varepsilon)| = \lim_{\varepsilon \to 0} \left|\frac{\gamma(\varepsilon)}{\nu(\varepsilon)}\right| = \infty,$$

and

$$\lim_{\varepsilon \to 0} |\mu(\varepsilon)| = \lim_{\varepsilon \to 0} \left|\frac{1}{\nu(\varepsilon)}\right| = \infty$$

since $\nu(\varepsilon) \to 0$ and $\gamma(\varepsilon) \to \gamma_* \neq 0$. We also have

$$\lim_{\varepsilon \to 0} \frac{\lambda(\varepsilon)}{\mu(\varepsilon)} = \lim_{\varepsilon \to 0} \gamma(\varepsilon) = \gamma_*.$$

The implication is thus proven.

To prove that statement 2 implies statement 1, we suppose that the functions $\lambda, \mu : \mathcal{D} \to \mathbb{C} \times \mathbb{C}$ exists, with the desired limits. Define $\gamma(\varepsilon) := \lambda(\varepsilon)/\mu(\varepsilon)$ and $\nu(\varepsilon) := 1/\mu(\varepsilon)$. By Lemma 3.7, $(\gamma(\varepsilon), \nu(\varepsilon)) \in \mathbb{D}(B, A, \varepsilon)$. Since $\lambda(\varepsilon) \in \sigma(A + \mu(\varepsilon)B)$ and $\lambda(\varepsilon) \to \infty$ we must also have $\mu(\varepsilon) \to \infty$ as $\varepsilon \to 0$.

Consequently,

$$\lim_{\varepsilon \to 0} \nu(\varepsilon) = \lim_{\varepsilon \to 0} \frac{1}{\mu(\varepsilon)} = 0,$$

and $(\gamma(\varepsilon), \nu(\varepsilon)) \to (\gamma_*, 0)$ as $\varepsilon \to 0$. By point *(ii)* of Theorem 3.5, γ_* is then a double eigenvalue of B. This concludes the proof. $\square$

This illustrates that there are situations where some elements of $\mathbb{D}(A, B, \varepsilon)$ do not converge to a (finite) value. However, from statement (iii) of Theorem 3.5, we have that elements which do not converge, also do not correspond to elements of $\mathbb{B}(A, B)$. This justifies that although elements of $\mathbb{D}(A, B, \varepsilon)$ are not always convergent, we still capture all solutions of $\mathbb{B}(A, B)$.

3.3. Regularization error. Since ε must be chosen non-zero and not too small in practice (see Remark 3.3), the construction with relative distance will always generate some error, which we will call the regularization error. In this section we will see that this error is reasonably behaved. We show this by proving some results about the asymptotic error as a function of the regularization parameter ε. In particular, if the repeated eigenvalue has a completely regular square root splitting, which is the generic case, the accuracies are $\lambda(\varepsilon) - \lambda_* = \mathcal{O}(\varepsilon)$, i.e., linear, and $\mu(\varepsilon) - \mu_* = \mathcal{O}(\varepsilon^2)$, i.e., quadratic.

We will later use this result to propose a rough argument for how to choose ε in practice.

THEOREM 3.9 (Regularization error). *Let (λ_*, μ_*) be an isolated pair in $\mathbb{B}(A, B)$, then there exists a pair of functions $(\lambda(\varepsilon), \mu(\varepsilon)) \in \mathbb{D}(A, B, \varepsilon)$ such that*

$$\lim_{\varepsilon \to 0} \frac{|\lambda(\varepsilon) - \lambda_*|}{|\varepsilon|^\alpha} = C_1 |\lambda_*|^\alpha, \tag{3.14}$$

and

$$\lim_{\varepsilon \to 0} \frac{|\mu(\varepsilon) - \mu_*|}{|\varepsilon|^\beta} = C_2 |\lambda_*|^\beta, \tag{3.15}$$

where $C_1, C_2 \in \mathbb{R}$ and $\alpha, \beta \in \mathbb{R}_+$ are independent of ε. Moreover, if $\lambda(\cdot)$ has a completely regular square root splitting with $\lambda(\varepsilon) = \lambda_* \pm C(\mu_* - \mu(\varepsilon))^{1/2} + o(\mu_* - \mu(\varepsilon))$ and $C \neq 0$, then the constants are $\alpha = 1$, $\beta = 2$, $C_1 = \frac{1}{2}$ and $C_2 = \frac{1}{|4C^2|}$.

Proof. The proof of (3.14) and (3.15) is done for the general setting where the eigenvalue $\lambda(\varepsilon)$ has a Puiseux expansion,

$$\lambda(\varepsilon) = \lambda_* + C(\mu(\varepsilon) - \mu_*)^q + o((\mu(\varepsilon) - \mu_*)^q), \tag{3.16}$$

and $\lambda_2(\varepsilon)$ is either a different branch from the same Puiseux series or a different Puiseux expansion. First consider the Puiseux expansion of the difference,

$$\lambda(\varepsilon) - \lambda_1(\varepsilon) = D(\mu(\varepsilon) - \mu_*)^p + o((\mu(\varepsilon) - \mu_*)^p),$$

where $C, D \in \mathbb{C}$. Note that C and D can be chosen non-zero since (λ_*, μ_*) is isolated. The two equations can now be solved by eliminating $(\mu(\varepsilon) - \mu_*)^{qp}$,

$$((\lambda(\varepsilon) - \lambda_*)^p / C^p + o(\mu(\varepsilon) - \mu_*)^q)^p = (\mu(\varepsilon) - \mu_*)^{qp} =$$
$$(\lambda(\varepsilon) - \lambda_2(\varepsilon))^q / D^q + o(\mu(\varepsilon) - \mu_*)^p)^q.$$

The relative distance is ε by construction, i.e., $\lambda_2(\varepsilon) = \lambda(\varepsilon)(1 + \varepsilon)$. Inserting $\lambda_2(\varepsilon) = \lambda(\varepsilon)(1 + \varepsilon)$ into the second equality, yields that

$$|\lambda(\varepsilon) - \lambda_*|^p / |C|^p = |\varepsilon|^q |\lambda(\varepsilon)|^q / |D|^q + o((\mu(\varepsilon) - \mu_*)^p).$$

Hence, by taking the pth root,

$$\frac{|\lambda(\varepsilon) - \lambda_*|}{|\varepsilon|^{q/p}} = \frac{|C|}{|D|^{q/p}} |\lambda(\varepsilon)|^{q/p} + o(\mu(\varepsilon) - \mu_*).$$

This proves (3.14). From (3.16) we find that

$$|\mu(\varepsilon) - \mu_*| = \frac{(|\lambda(\varepsilon)||\varepsilon|)^{1/p}}{|D|^{1/p}} + o(|\varepsilon|^{1/p}).$$

This implies (3.15). For a completely regular square root splitting we can choose $p = q = 1/2$ and $D = 2C$. Hence, $\alpha = 1$, $\beta = 1/p = 2$, $C_1 = |C/2C| = 1/2$, $C_2 = 1/(4C^2)$. $\square$

REMARK 3.10 (Extrapolation). *The algebraic convergence of the solution as $\varepsilon \to 0$ can be used to accelerate convergence by evaluating for several ε. This extrapolation technique will not be used in this work since evaluating $(\lambda(\varepsilon), \mu(\varepsilon))$ is the most computationally dominating part of the method. In order to gain full accuracy we will (in Section 4) instead use a local iterative method to gain high accuracy.*

3.4. The regularization parameter trade-off. In practice we must fix the regularization parameter ε before the carrying out the algorithm. If ε is chosen too large, the regularization error will (according to Theorem 3.9) be large. If ε is chosen too small it is natural to expect that rounding errors will destroy the accuracy of the solution since the problem is very close to singular. We now wish to derive a rough estimate of ε providing a reasonable resolution to the trade-off. We will restrict the study to the generic case where λ is a double (not triple) eigenvalue of $A + \mu B$.

In this subsection we will also restrict the error analysis to the matrix determinant approach to solve the two-parameter eigenvalue problem, i.e., we will solve the generalized eigenvalue problem corresponding to the pencil

$$D_\varepsilon(\lambda) := \Delta_1 - \lambda \Delta_0(\varepsilon).$$

We now wish to outline in what way the magnitude of ε affects the condition of the eigenvalue problem of D_ε for $\varepsilon \neq 0$. This will lead us to a reasonable choice of ε to be used in the implementation.

We start with a technical lemma showing that the structure of the problem is such that the first terms in the expansion of $\det(D_\varepsilon(\lambda))$ (in powers of ε) vanish.

LEMMA 3.11. *The following expansion holds:*

$$\det D_\varepsilon(\lambda) = \sum_{k=2}^{n^2} \varepsilon^k f_k(\lambda),$$

where the functions f_k, $k = 2, \ldots, n^2$, are polynomials of degree smaller than or equal to n^2.

Proof. By explicitly computing the determinant it follows that

$$\det D_\varepsilon(\lambda) = \sum_{k=0}^{n^2} \varepsilon^k f_k(\lambda),$$

where f_k, $k = 0, \ldots, n^2$, are polynomials. Due to the fact that the pencil $D_0(\lambda)$ is singular, we have $f_0 \equiv 0$. It remains to prove that $f_1 \equiv 0$.

For an arbitrary fixed value of λ, we get

$$f_1(\lambda) \quad = \frac{\partial \det D_\varepsilon(\lambda)}{\partial \varepsilon}\Big|_{\varepsilon=0} = \operatorname{Tr}\left\{\operatorname{adj}(D_0(\lambda)) \frac{\partial D_\varepsilon(\lambda)}{\partial \varepsilon}\Big|_{\varepsilon=0}\right\}.$$

We can express

$$\begin{aligned} D_0(\lambda) \quad &= (-A \otimes B + B \otimes A) - \lambda(-I \otimes B + B \otimes I) \\ &= B \otimes (A - \lambda I) - (A - \lambda I) \otimes B. \end{aligned}$$

If $(A - \lambda I)$ is invertible, then we get from Lemma A.1 in the appendix that $\operatorname{adj} D_0(\lambda) = 0$. We conclude that $f_1(\lambda) = 0$ for all $\lambda \notin \sigma(A)$. Because f_1 is a polynomial this implies that $f_1 \equiv 0$ and the proof is completed. □

In what follows we investigate the condition of the eigenvalue problem of the pencil D_ε, for a fixed value of $\varepsilon \neq 0$. In order to assess the effect of perturbations of the matrices Δ_0 and Δ_1 on the eigenvalues of D_ε we consider the pseudospectra $\Lambda_\gamma(D_\varepsilon)$, $\gamma > 0$, defined here as:

$$\Lambda_\gamma(D_\varepsilon) := \Big\{\lambda \in \mathbb{C}: \ \det\{(\Delta_1 + \delta\Delta_1) - \lambda(\Delta_0(\varepsilon) + \delta\Delta_0)\} = 0 \text{ for some}$$

$$\delta\Delta_0, \delta\Delta_1 \in \mathbb{C}^{n^2 \times n^2}, \text{ satisfying } \frac{\|\delta\Delta_0(\varepsilon)\|_2}{\|\Delta_0\|_2} < \gamma \text{ and } \frac{\|\delta\Delta_1\|_2}{\|\Delta_1\|_2} < \gamma\Big\}. \quad (3.17)$$

Thus, the pseudospectrum $\Lambda_\gamma(D_\varepsilon)$ is a subset of the complex plane consisting of all possible positions of the eigenvalues of D_ε when the system matrices are subject to perturbations with relative size smaller than γ. From [21, Theorem 1] the following computational formula can be derived:

$$\Lambda_\gamma(D_\varepsilon) = \Big\{\lambda \in \mathbb{C}: \ \|(D_\varepsilon(\lambda))^{-1}\|_2(\|\Delta_0(\varepsilon)\|_2 + |\lambda|\|\Delta_1\|_2) > \frac{1}{\delta}\Big\}. \quad (3.18)$$

Now, let λ_ε be an isolated eigenvalue of D_ε, that is,

$$\det D_\varepsilon(\lambda_\varepsilon) = 0, \quad \frac{d}{d\lambda}\det D_\varepsilon(\lambda)\Big|_{\lambda=\lambda_\varepsilon} \neq 0.$$

If $|\lambda - \lambda_\varepsilon|$ is small we can approximate

$$\|(D_\varepsilon(\lambda))^{-1}\|_2(\|\Delta_0(\varepsilon)\|_2 + |\lambda|\|\Delta_1\|_2)$$

$$\approx \frac{\|\operatorname{adj} D_\varepsilon(\lambda_\varepsilon)\|_2}{\left|\frac{d}{d\lambda}\det D_\varepsilon(\lambda)\big|_{\lambda=\lambda_\varepsilon}\right||(\lambda - \lambda_\varepsilon)|}(\|\Delta_0(\varepsilon)\|_2 + |\lambda_\varepsilon|\|\Delta_1\|_2). \quad (3.19)$$

From (3.18) and (3.19) we conclude that for sufficiently small values of γ, the pseudospectrum Λ_γ contains a disc centered around λ_ε, with radius equal to $r(\lambda_\varepsilon)\delta$, where

$$r(\lambda_\varepsilon) := \frac{\|\text{adj } D_\varepsilon(\lambda_\varepsilon)\|_2}{\left|\frac{d}{d\lambda} \det D_\varepsilon(\lambda)\big|_{\lambda=\lambda_\varepsilon}\right|} (\|\Delta_0(\varepsilon)\|_2 + |\lambda_\varepsilon| \|\Delta_1\|_2). \tag{3.20}$$

In words, the number $r(\lambda_\varepsilon)$ is the growth rate of the pseudospectrum $\Lambda_\gamma(D_\varepsilon)$ around the eigenvalue λ_ε when δ is increased from zero.

REMARK 3.12. *The number $r(\lambda_\varepsilon)$ corresponds to the (structured) condition number for the eigenvalue λ_ε as defined in [1, Equation (4)]. We will now use the property $\text{adj } D_\varepsilon(\lambda_\varepsilon) = c\, uv^*$, where u and v are normalized right and left null vectors of $D_\varepsilon(\lambda_\varepsilon)$ and $c \in \mathbb{C}$. Moreover,*

$$\frac{d}{d\lambda} \det D_\varepsilon(\lambda)\bigg|_{\lambda=\lambda_\varepsilon} = \text{tr}\left(D'_\varepsilon(\lambda_\varepsilon) \text{ adj } D_\varepsilon(\lambda_\varepsilon)\right) = cv^* D'_\varepsilon(\lambda_\varepsilon)u.$$

The expression (3.20) can now be simplified to

$$r(\lambda_\varepsilon) = \frac{1}{|v^* D'_\varepsilon(\lambda_\varepsilon)u|} (\|\Delta_0(\varepsilon)\|_2 + |\lambda_\varepsilon| \|\Delta_1\|_2),$$

which is consistent with the expression formulated in [1, Lemma 2.1].

Taking into account Lemma 3.11 we can simplify (3.20) to

$$r(\lambda_\varepsilon) = \frac{1}{\varepsilon^2} \frac{\|\text{adj } D_\varepsilon(\lambda_\varepsilon)\|_2}{\left|\sum_{k=2}^{n^2} f'_k(\lambda_\varepsilon)\varepsilon^{k-2}\right|} (\|\Delta_0(\varepsilon)\|_2 + |\lambda_\varepsilon| \|\Delta_1(\lambda_\varepsilon)\|_2).$$

Hence, if $\lim_{\varepsilon\to 0} \lambda_\varepsilon$ is finite, say λ_*, then the pseudospectral growth rate increases inversely proportional to ε^2 as $\varepsilon \to 0$. We will now reach the main point of this section. On the one hand, if we apply a stable algorithm to compute the eigenvalues of D_ε (for instance, the celebrated QZ algorithm) it is expected that the computational error on the result is comparable to the error induced by rounding errors on the data. Therefore, for a fixed value of ε the worst-case computational error on the eigenvalue λ_ε is expected to be proportional to

$$E_c(\varepsilon) := r(\lambda_\varepsilon) \, \varepsilon_{\text{mach}} \sim \frac{\varepsilon_{\text{mach}}}{\varepsilon^2}, \tag{3.21}$$

where $\varepsilon_{\text{mach}}$ is the machine precision. On the other hand, recall from Theorem 3.9 that the approximation error, that is $E_a := (\lambda(\varepsilon) - \lambda_*)/\lambda_*$, satisfies

$$E_a(\varepsilon) \sim \varepsilon, \tag{3.22}$$

which needs to be small, in order to obtain good approximations of all solutions of $\mathbb{B}(A, B)$. The optimal choice of ε involves a trade-off between (3.21) and (3.22). It leads us to the choice

$$\varepsilon \sim \varepsilon_{\text{mach}}^{1/3}, \tag{3.23}$$

for which both

$$E_c \sim \varepsilon_{\text{mach}}^{1/3}, \quad E_a \sim \varepsilon_{\text{mach}}^{1/3}.$$

As demonstrated in Section 5, this must be treated as a rule of thumb.

4. Local methods. The method in Section 3 (MFRD) has the attractive property that it approximates all solutions of $\mathbb{B}(A, B)$. A drawback of the method is that a regularization error is introduced dependent on the parameter ε which can not be chosen zero or too small. In this section, we show how, in a post-processing step, approximations can be improved using local iterative methods. These methods are based on solving systems of nonlinear equations that fully characterize the elements of $\mathbb{B}(A, B)$. Since they are iterative methods, they rely on having reasonable starting values. These can be generated from the solutions of $\mathbb{D}(A, B, \varepsilon)$, i.e., MFRD.

4.1. An augmented system of equations. A double eigenvalue can either be semisimple or non-semisimple. If it is semisimple then

$$(A + \mu B - \lambda I)v_1 = 0 \tag{4.1a}$$
$$(A + \mu B - \lambda I)v_2 = 0, \tag{4.1b}$$

for non-parallel v_1 and v_2, whereas for non-semisimple eigenvalues, we have a generalized eigenvector u associated with the eigenvector v, such that

$$(A + \mu B - \lambda I)v = 0 \tag{4.2a}$$
$$(A + \mu B - \lambda I)u = v. \tag{4.2b}$$

In our setting, we do not know a priori which case occurs. In order to construct a local iterative method which works for both semisimple and non-semisimple eigenvalues we will consider the null space of $(A + \mu B - \lambda I)^2$. It is easy to verify, by multiplying (4.1) and (4.2) from the left by $(A + \mu B - \lambda I)$, that $(A + \mu B - \lambda I)^2$ must have a null-space of at least dimension two, if λ is an eigenvalue of $A + \mu B$ of multiplicity two. The converse also holds [19, Lemma 1]. That is, we have that $A + \mu B$ has a multiple (semisimple or non-semisimple) eigenvalue λ if and only if $(A + \mu B - \lambda I)^2$ has a null-space of at least dimension two.

One common approach to the standard eigenvalue problem is to write it as a system of equations by introducing a normalization condition, e.g., $a^* v = 1$, where the normalization vector a can be arbitrary but not orthogonal to the eigenvector. We will use a similar construction, where we need several normalization constraints.

We will use the following general property of a matrix $E \in \mathbb{C}^{n \times n}$ with a two-dimensional null-space. For almost any $a_1, a_2 \in \mathbb{C}^n$, the set of equations

$$Ev_1 = 0 \tag{4.3a}$$
$$Ev_2 = 0 \tag{4.3b}$$
$$a_1^* v_1 = 1 \tag{4.3c}$$
$$a_2^* v_1 = 1 \tag{4.3d}$$
$$a_1^* v_2 = 1 \tag{4.3e}$$
$$v_1^* v_2 = 0 \tag{4.3f}$$

uniquely defines an orthogonal basis $(v_1, v_2) \in \mathbb{C}^{n \times 2}$ of the null-space of E.

This can be directly derived from the singular value decomposition of E. Let $W \in \mathbb{C}^{n \times (n-2)}$ denote the (right) singular vectors corresponding to the $n-2$ non-zero singular values of E. The singular vectors are orthonormal and hence $Ev_1 = 0$ is equivalent to $W^* v_1 = 0$. The equations (4.3a), (4.3c), (4.3d) can now be written as a

linear system

$$\begin{bmatrix} W^* \\ a_1^* \\ a_2^* \end{bmatrix} v_1 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix},$$

which has a unique solution v_1, provided that

$$\mathrm{rank} \begin{bmatrix} W^* \\ a_1^* \\ a_2^* \end{bmatrix} = n. \tag{4.4}$$

From the same reasoning with (4.3b), (4.3e) and (4.3f), we have that v_2 is unique, provided that

$$\mathrm{rank} \begin{bmatrix} W^* \\ a_1^* \\ v_1^* \end{bmatrix} = n. \tag{4.5}$$

We will now combine this result with [19, Lemma 1], stating that $E = (A + B\mu - \lambda I)^2$ has a null-space of dimension at least two, iff λ is a non-simple eigenvalue of $A + \mu B$. This leads us to the following equations, which characterize a multiple eigenvalue λ of $A + \mu B$,

$$(A + B\mu - \lambda I)^2 v_1 = 0 \tag{4.6a}$$
$$(A + B\mu - \lambda I)^2 v_2 = 0 \tag{4.6b}$$
$$a_1^* v_1 = 1 \tag{4.6c}$$
$$a_2^* v_1 = 1 \tag{4.6d}$$
$$a_1^* v_2 = 1 \tag{4.6e}$$
$$v_1^* v_2 = 0 \tag{4.6f}$$

for any a_1 and a_2 satisfying (4.4) and (4.5), i.e., for almost any choice of a_1 and a_2.

4.2. A zero-residual Gauss-Newton iteration. If we consider (λ, μ, v_1, v_2) as unknowns, the system (4.6) has $2(n+1)$ unknowns and consists of $2(n+2)$ conditions. It is hence an over-determined system. The Gauss-Newton (GN) method is the standard approach to solve over-determined systems of equations, and we propose to use GN to solve (4.6). We briefly summarize GN for our setting. The Gauss-Newton method is a method to find the minimum of a residual in the two-norm sense. Given a function $r : \mathbb{R}^p \to \mathbb{R}^q$, which will correspond the difference between the left- and right-hand sides of (4.6), we wish to find the minimum of

$$f(x) = \frac{1}{2} r(x)^T r(x). \tag{4.7}$$

By denoting the Jacobian of r by $J : \mathbb{R}^p \to \mathbb{R}^{q \times p}$, the GN iteration can be seen as an iteration of the following steps.
- For the kth iterate $x^{(k)}$, compute the update $\Delta x^{(k)}$ as the least squares solution to the overdetermined system

$$r(x^{(k)}) + J(x^{(k)})\Delta x^{(k)} = 0. \tag{4.8}$$

16

- Update the iterate $x^{(k+1)} = x^{(k)} + \Delta x^{(k)}$.

See, e.g., [25, Section 10.2] for further details about GN including termination criteria.

We will use a relation with Newton's method in order to illustrate and predict the convergence of the method. The important property we will use is that (4.6) has an exact solution, and the minimum of $f(x)$ is also a zero of $r(x)$. In this situation (called the zero-residual case) GN has some nice convergence properties. In particular, the convergence is very similar to the convergence of Newton's method.

THEOREM 4.1. *Let $\Delta \tilde{x}^{(k)}$ be one step of the (standard) Newton method for the minimization of the function* (4.7). *Suppose $f(x_*) = 0$ and that r is sufficiently smooth. Then,*

$$\left(J(x^{(k)})^T J(x^{(k)}) + R_k\right) \Delta \tilde{x}^{(k)} = -J(x^{(k)})^T r(x^{(k)}), \tag{4.9}$$

where

$$R_k = O(x^{(k)} - x_*).$$

Proof. The Newton iteration to minimize $f(x)$ can be written as

$$H_k \Delta \tilde{x}^{(k)} = -J(x^{(k)})^T r(x^{(k)}), \tag{4.10}$$

where H_k is the Hessian of f in $x^{(k)}$. A straightforward computation (see also [25, Equation (10.5)]) yields

$$H_k = J(x^{(k)})^T J(x^{(k)}) + R_k, \tag{4.11}$$

where the element in position (i, j) of R_k is given by

$$(R_k)_{i,j} = \sum_{l=1}^{q} r_l(x^{(k)}) \frac{\partial^2 r_l(x^{(k)})}{\partial x_i x_j}, \ i, j = 1, \ldots, p.$$

In the zero-residual case, the matrix R_k vanishes in the solution x_* and from the smoothness of r we can write $R_k = O(x^{(k)} - x_*)$. $\square$

Now note that the correction $\Delta x^{(k)}$ in a GN step fulfills the normal equations, i.e., the least-squares solution to (4.8) satisfies

$$J(x^{(k)})^T J(x^{(k)}) \Delta x^{(k)} = -J(x^{(k)})^T r(x^{(k)}). \tag{4.12}$$

By comparing (4.12) and (4.9) we see that one step of GN is equal to one step of Newton's method in an asymptotic sense. Due to this result, we expect that the convergence of the Gauss-Newton method is similar to that of Newton's method. This is consistent with the following results which are partially formalized in [25, Section 10.2].

- If $J(x_*)^T J(x_*)$ is nonsingular, i.e., $J(x_*)$ has full column rank. Then, the Gauss-Newton method applied to r exhibits *quadratic convergence* to $x = x_*$.
- If $J(x_*)^T J(x_*)$ is singular, i.e., $J(x_*)$ is rank-deficient. Then, the Gauss-Newton method applied to r is expected to generically exhibit *linear convergence* to $x = x_*$.

We now turn to our specific case, i.e., the application of GN to solve the over-determined system (4.6). The rank of the Jacobian at the solution for this set of equations is discussed in Appendix B, where we find that the generic situation for non-semisimple eigenvalues is that the Jacobian has full rank. For semisimple eigenvalues the Jacobian is rank-deficient and we observe linear convergence in examples.

It remains to propose choices of initial values (for the vectors v_1 and v_2) and constant vectors a_1 and a_2. The characterization (4.6) is based on the assumption that a_1 and a_2 are given and fulfill the generic conditions (4.4) and (4.5). It is, in fact, possible to construct an appropriate choice for a_1 and a_2 from a sufficiently good approximate solution. Consider an approximation generated by MFRD, $(\tilde{\lambda}, \tilde{\mu}) \in \mathbb{D}(A, B, \varepsilon)$, for a small value of ε. Then an approximate basis $(\tilde{v}_1, \tilde{v}_2)$ for the two-dimensional nullspace of $(A + B\mu - \lambda I)^2$, with the property

$$\tilde{v}_1^* \tilde{v}_1 = 1, \ \tilde{v}_2^* \tilde{v}_2 = 1, \ \tilde{v}_1^* \tilde{v}_2 = 0,$$

can be generated from the singular value decomposition of $(A + B\tilde{\mu} - \tilde{\lambda}I)^2$. In the light of this, an appropriate choice for a_1 and a_2 in (4.6) is given by

$$a_1 = \tilde{v}_1 + \tilde{v}_2, \quad a_2 = \tilde{v}_1 - \tilde{v}_2. \tag{4.13}$$

In conclusion, a local correction can be performed by solving equations (4.6), with, e.g., the choice (4.13) for the normalization constraints and with the starting values $(\tilde{\lambda}, \tilde{\mu}, \tilde{v}_1, \tilde{v}_2)$.

REMARK 4.2 (Semi-simple). *In the case where the double eigenvalue is semi-simple, the Jacobian is rank-deficient. As mentioned above, a loss of quadratic convergence of GN is observed in the examples. As outlined also in Appendix B, this inconvenience can be overcome by instead applying GN to the equations*

$$(A + B\mu - \lambda I)v_1 = 0 \tag{4.14a}$$
$$(A + B\mu - \lambda I)v_2 = 0 \tag{4.14b}$$
$$a_1^* v_1 = 1 \tag{4.14c}$$
$$a_2^* v_1 = 1 \tag{4.14d}$$
$$a_1^* v_2 = 1 \tag{4.14e}$$
$$v_1^* v_2 = 0 \tag{4.14f}$$

which directly characterize the presence of a two-dimensional eigenspace of the matrix $A + B\mu$.

REMARK 4.3 (Relation with Newton's method). *In this paper we propose to use an iterative correction method based on solving an over-determined system of equations. It is possible to characterize a double eigenvalue with a fully determined system of equations. Consider the equations (4.6a), (4.6c) and (4.6d). With these conditions we require that the vector v_1 satisfies two normalization constraints simultaneously. A situation where this occurs is when v_1 belongs to a two-dimensional subspace, here, the null-space of $(A + \mu B - \lambda I)^2$. Since the (4.6a), (4.6c), (4.6d) do not involve v_2, we have $n + 2$ equations and $n + 2$ unknowns. It is a fully determined system and one could apply Newton's method to this set of equations. In this paper we do not suggest to use such an approach, but instead propose to use the overdetermined system (4.6) since GN applied to (4.6) appears to have several favorable numerical properties. The system (4.6) is a characterization of a two-dimensional null-space using two vectors which are orthogonal, whereas the fully determined system (4.6a), (4.6c), (4.6d) is a*

characterization using only one vector. These intuitive arguments and our numerical experiments serve as motivation for us to propose to solve the overdetermined system instead of the fully determined system.

REMARK 4.4. *In the case where λ_* is a* non-semisimple *eigenvalue of $A + \mu_* B$, an alternative method consists of applying the algorithm of [18]. The latter is based on an application of Newton's method to solve a system of equations that characterizes the presence of a Jordan block of at least two-by-two in the canonical representation. In our experiments it yields a performance comparable to solving (4.6) with Gauss-Newton.*
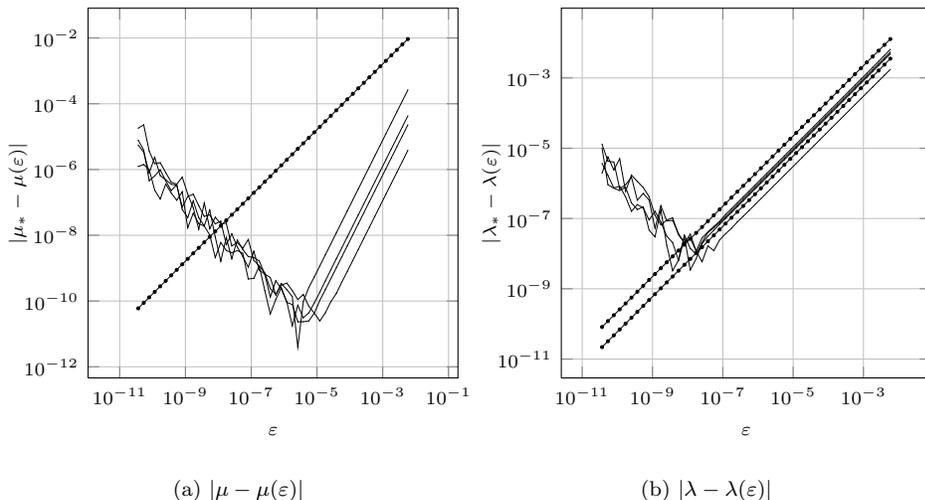

(a) $|\mu - \mu(\varepsilon)|$ (b) $|\lambda - \lambda(\varepsilon)|$

FIGURE 5.1. *Logarithmic plot of the accuracy in the MFRD approach as a function of ε for the example in Section 5.1. The figures show the error in μ_* (left) and λ_* (right), respectively. The lines with markers are the errors for the semisimple point (μ_0, λ_0), while the other lines are the non-semisimple points.*

5. Examples.

5.1. A 3×3 test problem. Suppose

$$A = \begin{bmatrix} -1 & 2 & 1 \\ 0 & 2 & -\mathrm{i} \\ \mathrm{i} & 1 & -\mathrm{i} \end{bmatrix}, \quad B = (C - A)/\mu_0, \quad C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}, \qquad (5.1)$$

where $\mu_0 = 1 + \mathrm{i}$. This problem is constructed such that for $\mu = \mu_0$, $\lambda = 2$ is a semisimple eigenvalue of $A + \mu B$. It can be solved explicitly with software for symbolic manipulations by simultaneously solving $f_\lambda(\lambda, \mu) = 0$ and $f(\lambda, \mu) = 0$ as in Remark 2.2. The solution is $\mathbb{B}(A, B) \approx \{(1 + i, 2), (0.60 + 0.40i, 0.50 - 0.39i), (0.98 + 1.4i, 1.6 + 0.32), (1.1 + 1.3i, 2.1 + 0.21i), (1.5 + 1.2i, 1.9 - 0.2i)\}$. We will use this solution (with sufficiently high precision) for reference.

This example will now be used to illustrate how the accuracy of the solution computed with MFRD depends on the regularization parameter ε. We will also illustrate the predicted behavior of the local iterative method of Section 4. The error of MFRD for the solutions is given in Figure 5.1. We observe V-shaped error curves (as predicted in Section 3.4) for the non-semisimple eigenvalues, corresponding essentially to

the maximum of the rounding error (dominating for small ε) and regularization error (dominating for larger ε).

Note also that there is no optimal choice of ε in the sense that the errors of the individual approximations will never be simultaneously minimized. This holds in particular for the semisimple eigenvalue for which the regularization error is of order $|\mu - \mu(\varepsilon)| = O(\varepsilon)$ (unlike the other error curves, for which $|\mu - \mu(\varepsilon)| = O(\varepsilon^2)$). Hence, there is no choice of ε such that MFRD returns full precision solutions. This problem is not present when we combine MFRD with the individual treatment of the approximations with the iterative method of Section 4.

The convergence of the iterative method is illustrated in Figure 5.2. For illustrative purposes we used $\varepsilon = 0.01$ instead of the estimate of a good choice given in Section 3.4. We clearly see that the convergence is superlinear for the non-semisimple eigenvalues and linear for the semisimple. Note that we can achieve essentially full precision for all eigenvalues, although it requires more iterations for the semisimple eigenvalue. We also observe (in Figure 5.2b), that the quadratic convergence for the semisimple eigenvalue can, as predicted, be restored by instead solving the equation (4.14).
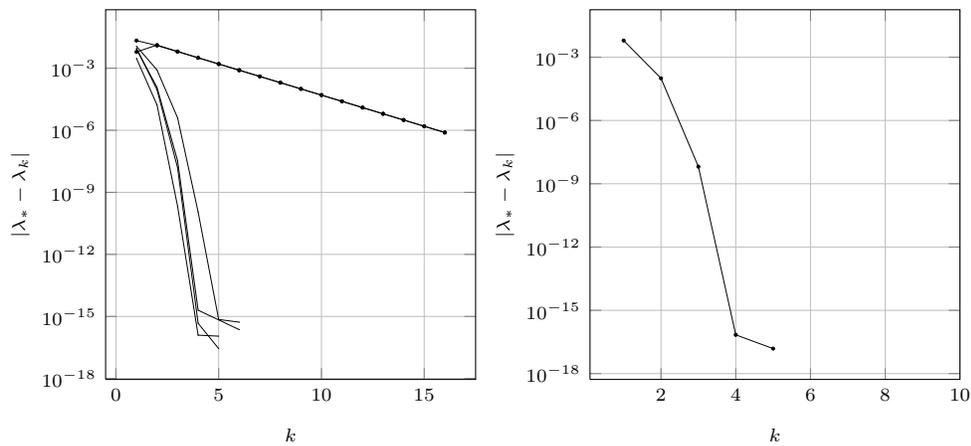
(a) Quadratic convergence is observed for the non-semisimple case and linear convergence for the semisimple case.

(b) By considering the system (4.14) the quadratic convergence for the semisimple eigenvalue can be recovered

FIGURE 5.2. *Illustration of local iterative methods applied to the example in Section 5.1.*

5.2. Application to quantum mechanical perturbation theory. An important problem in computational quantum mechanics is to characterize the smallest eigenvalues, assuming these exist, of a self-adjoint partial differential operator $\mathcal{H} = \mathcal{A} + \mu_{\mathrm{phys}}\mathcal{B}$ (with $\mu_{\mathrm{phys}} \in \mathbb{R}$) over some infinite dimensional Hilbert space. Standard discretization techniques produce a (Hermitean) matrix approximation $H = A + \mu_{\mathrm{phys}}B$ to this operator. The matrix dimensions may be very large, and direct computation of the eigenvalues may be infeasible. For the simplest case, where the smallest eigenvalue $\lambda_1(\mu)$ is to be determined, one considers the pencil $A + \mu B$ and

20

forms a Taylor series expansion

$$\lambda_1(\mu) = \sum_{k=0}^{\infty} c_k \mu^k, \qquad (5.2)$$

which is to be evaluated at $\mu = \mu_{\text{phys}}$. Clearly, $c_0 = \lambda_1(0)$, which is the smallest eigenvalue of A (assumed to be simple), and $c_1 = u^* B u$, with u being a normalized eigenvector of A corresponding to $\lambda_1(0)$ [16, Chapter II, Remark 2.2]. The expressions for c_k, $k > 1$, become increasingly complicated, and they are typically evaluated using diagrammatic rules.

The eigenvalue functions $\lambda_j(\mu)$ are algebraic functions having only branch point type singularities, which is generic for non-semisimple eigenvalues, at complex values of μ. These may occur only when $\lambda_j(\mu_*) = \lambda_k(\mu_*) = \lambda_*$ for $j \neq k$ [27]. It can be shown that the (in general unknown) radius of convergence of (5.2) is given by

$$R = \min\left\{|\mu_*| \ : \ \lambda_1(\mu_*) = \lambda_k(\mu_*), \ k > 1\right\}. \qquad (5.3)$$

We note that $R \geq \min\{|\mu_*|\}$, i.e., the bifurcation point μ_* that actually limits R is not necessarily the smallest. However, it is usually among the few smallest points.

For the more complicated problem of determining several of the smallest eigenvalues $\lambda_j(\mu)$ via so-called partitioning techniques [27] the limiting singularity is also among the branch points, but usually of much larger magnitude.

We visualize the set $\mathbb{B}(A, B)$ in Figure 5.3 for a simple example of dimension $n = 15$. The branch point limiting the radius of convergence is highlighted. (In this case $R = \min\{|\mu_*|\}$, and $\lambda_* = \lambda_1 = \lambda_2$.) The example is that of two electrons in a harmonic oscillator trap with the discretization described in [29]. The matrix A is diagonal with equally spaced eigenvalues $1, 3, 5, \cdots$, while B is dense with elements that decay algebraically, i.e., $B_{ij} = O[(ij)^{-\beta}]$, $\beta > 0$. The matrices thus have a high degree of structure clearly reflected in Figure 5.3.
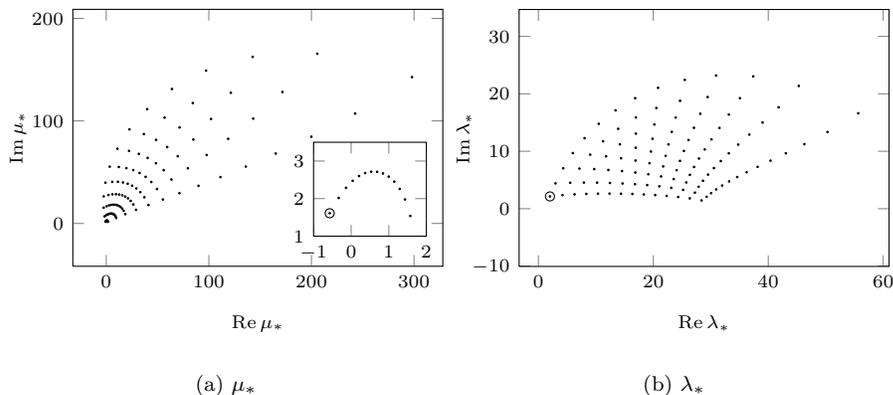

(a) μ_* (b) λ_*

FIGURE 5.3. *The set $\mathbb{B}(A, B)$ for a discretized partial differential operator from quantum mechanics. The left plot shows μ_* with the inset figure showing a magnification of the points closest to 0, while the right plot shows λ_*. The radius of convergence of the perturbation series is limited by the magnitude of the marked point in the μ_* plot, for which $\lambda_1(\mu_*) = \lambda_2(\mu_*) = \lambda_*$, the marked point in the λ_* plot.*

6. Conclusions and outlook. An important and unusual property of the method in this paper is the *global convergence* to all solutions, which can be combined with an iterative method to get full accuracy. The global convergence property comes from the observation that the fixed relative distance problem is a two-parameter eigenvalue problem.

The method can be adapted in many ways. Note that the elements of $\mathbb{B}(A, B)$ are characterized by a GEP in Section 3. There are many methods to solve GEPs and the right choice of method depends on the application. If only small number of solutions are of interest in the application, e.g., those solutions where $|\mu_*|$ is small, we can use an eigenvalue solver which focuses on these solutions. As an alternative promising approach, one can solve the two-parameter eigenvalue problem directly with methods which allow focusing on a target, e.g., [11].

We conclude this paper with some straightforward extensions following from the connection with the two-parameter eigenvalue problem. The method of fixed relative distance can be adapted to generalized eigenvalue problems. That is, the problem of finding (λ, μ) such that $\lambda C x = (A + \mu B)x$, $x \in \mathbb{C}^n \backslash \{0\}$, where λ is a double eigenvalue, can also be solved by considering the fixed relative distance and solving a two-parameter eigenvalue problem.

The multiparameter eigenvalue problem is a generalization of the two-parameter eigenvalue problem. Consider the problem of finding $\lambda, \mu_1, \mu_2, \ldots, \mu_k$ such that λ is an eigenvalue of multiplicity $k + 1$ of the eigenvalue matrix of $A + B_1\mu_1 + \cdots + B_k\mu_k$. This problem can also be solved with an approach based on fixed relative distance. By considering a relative disturbance of each of the parameters $\mu_1, \ldots, \mu_k$ we have a system of equations which can be interpreted as a $(k + 1)$-parameter eigenvalue problem.

Appendix A. A technical lemma.
LEMMA A.1. *For all $U, V \in \mathbb{C}^{n \times n}$, where either U or V is invertible, we have*

$$\text{adj}\,(U \otimes V - V \otimes U) = 0. \tag{A.1}$$

Proof. Because the role of U and V can be interchanged we assume, without loosing generality, that V is nonsingular. We further assume that $n \geq 2$, since the result is trivial for $n = 1$.

First, we characterize the null space of

$$U \otimes V - V \otimes U. \tag{A.2}$$

From the invertibility of V we have

$$(U \otimes V - V \otimes U)X = 0, \quad X \in \mathbb{C}^{n^2 \times 1},$$

if and only if

$$((V^{-1}U) \otimes I - I \otimes (V^{-1}U))X = 0,$$

22

i.e.,

$$((V^{-1}U) \oplus (-V^{-1}U))X = 0. \tag{A.3}$$

If $V^{-1}U$ has (at least) two linearly independent eigenvectors, E_1 and E_2, then it follows by inspection that $X = E_1 \otimes E_1$ and $X = E_2 \otimes E_2$ are linearly independent solutions of (A.3). In the other case, there always exists an eigenvalue λ, an eigenvector E and a generalized eigenvector H satisfying

$$(V^{-1}U)E = \lambda E, \ (V^{-1}U)H = \lambda H + E,$$

and it can be verified that $X = E \otimes E$ and

$$X = E \otimes H + H \otimes E$$

are linearly independent solutions of (A.3). Hence, in both cases we conclude that the null space of (A.2) has at least dimension two. This implies that there exists a matrix $T \in \mathbb{C}^{n^2 \times n^2}$ such that

$$T(U \otimes V - V \otimes U)T^{-1} = \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix},$$

where R_{22} is the two-by-two null matrix. We get

$$\mathrm{adj}(T^{-1}) \, \mathrm{adj}(U \otimes V - V \otimes U) \, \mathrm{adj}(T) = \mathrm{adj}\left(\begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix} \right) = 0.$$

The statement of the lemma follows. $\square$

Appendix B. The column rank of the Jacobian of (4.6) **and** (4.14)**.** The Jacobian matrix of (4.6) in the solution, given by

$$J := \begin{bmatrix} M^2 & 0 & (MB+BM)v_1 & -2Mv_1 \\ 0 & M^2 & (MB+BM)v_2 & -2Mv_2 \\ a_1^* & 0 & 0 & 0 \\ a_2^* & 0 & 0 & 0 \\ 0 & a_1^* & 0 & 0 \\ v_2^T & v_1^* & 0 & 0 \end{bmatrix}, \tag{B.1}$$

where $M = \lambda I - A - \mu B$. The column rank of the Jacobian will be studied for two important cases.

Case 1: λ is a double, non-semisimple eigenvalue of $A+B\mu$.

Since a double non-semisimple eigenvalue corresponds to a Jordan block of dimension two, we know that there exists a matrix T such that

$$\tilde{M} := T^{-1}MT = \left[\begin{array}{cc|c} 0 & 1 & \\ 0 & 0 & \\ \hline & & R \end{array} \right],$$

where $R \in \mathbb{C}^{(n-2)\times(n-2)}$ is invertible. By pre-multiplying J with $\mathrm{diag}(T^{-1}, T^{-1}, 1, 1, 1, 1)$ and post-multiplying with $\mathrm{diag}(T, T, 1, 1)$ we obtain

$$\tilde{J} := \begin{bmatrix} \tilde{M}^2 & 0 & (\tilde{M}\tilde{B}+\tilde{B}\tilde{M})\tilde{v}_1 & -2\tilde{M}\tilde{v}_1 \\ 0 & \tilde{M}^2 & (\tilde{M}\tilde{B}+\tilde{B}\tilde{M})\tilde{v}_2 & 0 \\ \tilde{a}_1^* & 0 & 0 & 0 \\ \tilde{a}_2^* & 0 & 0 & 0 \\ 0 & \tilde{a}_1^* & 0 & 0 \\ w_2^T & w_1^* & 0 & 0 \end{bmatrix},$$

23

where

$$\tilde{B} = T^{-1}BT, \ \tilde{v}_1 = T^{-1}v_1, \ \tilde{a}_1^* = a_1^*T, \ \tilde{a}_2^* = a_2^*T, \ w_1^* = v_1^*T, \ w_2^T = v_2^*T.$$

Note that

$$\tilde{M}^2 = \left[\begin{array}{cc|c} 0 & 0 & \\ 0 & 0 & \\ \hline & & R^2 \end{array} \right]. \tag{B.2}$$

It is easy to see that with a random choice of vectors a_1 and a_2 in the normalization constraints in (4.6), the matrix obtained by taking the first two block columns of $\tilde{J}$ has full column rank with probability one, and, in addition, $\tilde{M}\tilde{v}_1 \neq 0$. Therefore, taking into account the structure of (B.2), the matrix $\tilde{J}$ (and thus J) has full column rank if the 2-by-2 matrix obtained by considering the first two rows of

$$[(\tilde{M}\tilde{B} + \tilde{B}\tilde{M})\tilde{v}_1 \ -2\tilde{M}\tilde{v}_1]$$

is invertible. This is the case (with probability one for a random choice of a_1 and a_2) if and only if the element at position (2,1) is nonzero, that is,

$$e_2^T(\tilde{M}\tilde{B} + \tilde{B}\tilde{M})\tilde{v}_1 \neq 0, \tag{B.3}$$

where $e_2 = [0\ 1\ 0\cdots\ 0]^T$ is a unity vector in $\mathbb{C}^n$. Considering that $e_2^T\tilde{M} = 0$ the condition (B.3) becomes, in terms of the original matrices,

$$U_0^*BV_0 \neq 0, \tag{B.4}$$

where $U_0 := T^{-*}e_2$ is the left null vector of M and $V_0 := Mv_1$ the right null vector of M.

The condition (B.4) can be rephrased as

$$\frac{\partial}{\partial\mu}\det(\lambda I - A - \mu B) \neq 0,$$

which is a necessary and sufficient condition for the complete regular splitting property of the eigenvalue λ of $A + \mu B$, see [14].

Recapitulating the above results, we arrive at the following proposition.

PROPOSITION B.1. *Let $(\lambda, \mu) \in \mathbb{B}(A, B)$ be such that λ is a double, non-semisimple eigenvalue of $A + \mu B$ satisfying the completely regular splitting property. Then the Jacobian matrix of (4.6) corresponding to (λ, μ) generically has full column rank.*

We should note here that the value of the condition number of the Jacobian depends on the choice of the vectors a_1 and a_2 in the normalization.

Case 2: λ is a double semisimple eigenvalue of $A + \mu B$.

The Jacobian (B.1) cannot be of full column rank in the solution of (4.6) because $Mv_1 = 0$. However, using the same arguments as used in the case above, it can be shown that the Jacobian of (4.14) in the corresponding solution generically has full column rank.

REFERENCES

[1] B. Adhikari, R. Alam, and D. Kressner. Structured eigenvalue condition numbers and linearizations for matrix polynomials. Technical report, ETH Zürich, 2008.

[2] R. Alam and S. Bora. On sensitivity of eigenvalues and eigendecompositions of matrices. *Linear Algebra Appl.*, 396:273–301, 2005.

[3] F. Atkinson. Multiparameter spectral theory. *Bull. Am. Math. Soc.*, 74:1–27, 1968.

[4] F. Atkinson. *Multiparameter Eigenvalue Problems. Volume I: Matrices and compact operators*. Academic press, New York, 1972.

[5] H. Baumgärtel. *Analytic perturbation theory for matrices and operators*. Operator Theory: Advances and applications Vol 15. Birkhäuser Verlag, 1985.

[6] J. V. Burke, A. S. Lewis, and M. Overton. Optimization and pseudospectra, with applications to robust stability. *SIAM J. Matrix Anal. Appl.*, 25(1):80–104, 2003.

[7] F. Chaitin-Chatelin, A. Harrabi, and A. Ilahi. About Hölder condition numbers and the stratification diagram for defective eigenvalues. *Math. Comput. Simul.*, 54:397–402, 2000.

[8] N. Cottin. Dynamic model updating: A multiparameter eigenvalue problem. *Mechanical Systems and Signal Processing*, 15(4):649–665, 2001.

[9] G. Golub and J. Wilkinson. Ill-conditioned eigensystems and the computation of the Jordan canonical form. *SIAM Rev.*, 18:578–619, 1976.

[10] D. Hinrichsen and A. J. Pritchard. *Mathematical Systems Theory I*. Springer Berlin, 2005.

[11] M. E. Hochstenbach, T. Košir, and B. Plestenjak. A Jacobi–Davidson type method for the two-parameter eigenvalue problem. *SIAM J. Matrix Anal. Appl.*, 26(2):477–497, 2005.

[12] M. E. Hochstenbach and B. Plestenjak. Backward error, condition numbers, and pseudospectra for the multiparameter eigenvalue problem. *Linear Algebra Appl.*, 375:63–81, 2003.

[13] R. Horn and C. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1991.

[14] R. Hryniv and P. Lancaster. On the perturbation of analytic matrix functions. *Integral Equations Oper. Theory*, 34(3):325–338, 1999.

[15] M. Karow. Eigenvalue condition numbers and a formula of Burke, Lewis and Overton. *Electron. J. Linear Algebra*, 15:143–153, 2006.

[16] T. Kato. *Perturbation theory for linear operators*. Springer-Verlag, 1995.

[17] D. Kressner, M. J. Peláez, and J. Moro. Structured Hölder condition numbers for multiple eigenvalues. *SIAM J. Matrix Anal. Appl.*, 31(1):175–201, 2009.

[18] A. A. Mailybaev. Computation of multiple eigenvalues and generalized eigenvectors for matrices dependent on parameters. *Numer. Linear Algebra Appl.*, 13:419–436, 2006.

[19] A. N. Malyshev. A formula for the 2-norm distance from a matrix to the set of matrices with multiple eigenvalues. *Numer. Math.*, 83(3):443–454, 1999.

[20] E. Mengi. Locating a nearest matrix with an eigenvalue of prescribed algebraic multiplicity. *Numer. Math.*, 2010. available online, DOI: 10.1007/s00211-010-0326-3O.

[21] W. Michiels, K. Green, T. Wagenknecht, and S.-I. Niculescu. Pseudospectra and stability radii for analytic matrix functions with application to time-delay systems. *Lin. Alg. Appl.*, 418(1):315–335., 2006.

[22] J. Moro, J. V. Burke, and M. Overton. On the Lidskii-Vishik-Lyusternik perturbation theory for eigenvalues of matrices with arbitrary jordan structure. *SIAM J. Matrix Anal. Appl.*, 18(4):793–817, 1997.

[23] A. Muhič and B. Plestenjak. On the quadratic two-parameter eigenvalue problem and its linearization. Technical report, University of Ljubljana, 2009.

[24] A. Muhič and B. Plestenjak. On the singular two-parameter eigenvalue problem. *Electron. J. Linear Algebra*, 18:420–437, 2009.

[25] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer-Verlag New York, 1999.

[26] A. Ruhe. Properties of a matrix with a very ill-conditioned eigenproblem. *Numer. Math.*, 15:57–60, 1970.

[27] T. H. Schucan and H. A. Weidenmüller. Perturbation theory for the effective interaction in nuclei. *Ann. Physics*, 76:483–501, 1973.

[28] J. Sun. On condition numbers of a nondefective multiple eigenvalue. *Numer. Math.*, 61(2):265–275, 1992.

[29] A. Wensauer, M. Korkusinski, and P. Hawrylak. Configuration interaction method for Fock-Darwin states. *Solid State Comm.*, 130:115–120, 2004.