Fast Computation of $\operatorname{Sep}_{\lambda}$ via Interpolation-based Globality Certificates

Tim Mitchell*

November 12, 2019

Revised: September 13, 2020, May 25, 2021, March 22, 2022, January 21, 2023

Abstract

Given two square matrices A and B, we propose a new approach for computing the smallest value $\varepsilon \geq 0$ such that A+E and A+F share an eigenvalue, where $||E|| = ||F|| = \varepsilon$. In 2006, Gu and Overton proposed the first algorithm for computing this quantity, called sep_{λ}(A, B) ("sep-lambda"), using ideas inspired from an earlier algorithm of Gu for computing the distance to uncontrollability. However, the algorithm of Gu and Overton is extremely expensive, which limits it to the tiniest of problems, and until now, no other algorithms have been known. Our new algorithm can be orders of magnitude faster and can solve problems where A and B are of moderate size. Moreover, our method consists of many "embarrassingly parallel" computations, and so it can be further accelerated on multi-core hardware. Finally, we also propose the first algorithm to compute an earlier version of sep-lambda where $||E|| + ||F|| = \varepsilon$.

Keywords: sep-lambda, eigenvalue separation, eigenvalue perturbation, pseudospectra, Hamiltonian matrix

Notation: $\|\cdot\|$ denotes the spectral norm, $\sigma_{\min}(\cdot)$ the smallest singular value, $\Lambda(\cdot)$ the spectrum, $\kappa(\cdot)$ the condition number of a matrix with respect to the spectral norm, $J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}$, a matrix $A \in \mathbb{C}^{2n \times 2n}$ is Hamiltonian if $(JA)^* = JA$, $\mu(\cdot)$ the Lebesque measure on \mathbb{R} , and bd \mathcal{A} , int \mathcal{A} , and cl \mathcal{A} respectively the boundary, interior, and closure of a set \mathcal{A} .

1 Introduction

The quantity $\operatorname{sep}_{\lambda}(A, B)$ measures how close two square matrices $A \in \mathbb{C}^{m \times m}$ and $B \in \mathbb{C}^{n \times n}$ are to sharing a common eigenvalue, in the sense of how much A and B must be perturbed in order to make this so. Varah first introduced $\operatorname{sep}_{\lambda}(A, B)$ in 1979 in [Var79], and it was subsequently studied by Demmel in [Dem83, Dem86, Dem87], although Demmel used a slightly modified version, partly "because it lets us state slightly sharper results later on" [Dem83, p. 24]. The two definitions are:

$$\operatorname{sep}_{\lambda}^{\mathcal{V}}(A,B) \coloneqq \min_{\substack{E \in \mathbb{C}^{m \times m} \\ F \in \mathbb{C}^{n \times n}}} \{ \varepsilon : \Lambda(A+E) \cap \Lambda(B+F) \neq \emptyset, \|E\| + \|F\| \le \varepsilon \},$$
(1.1a)

$$\operatorname{sep}_{\lambda}^{\mathcal{D}}(A,B) \coloneqq \min_{\substack{E \in \mathbb{C}^{m \times m} \\ F \in \mathbb{C}^{n \times n}}} \{ \varepsilon : \Lambda(A+E) \cap \Lambda(B+F) \neq \varnothing, \max(\|E\|, \|F\|) \le \varepsilon \},$$
(1.1b)

with $\operatorname{sep}_{\lambda}^{\mathcal{V}}(A, B)$ denoting Varah's definition and $\operatorname{sep}_{\lambda}^{\mathcal{D}}(A, B)$ denoting Demmel's. Obviously, they are both zero if A and B share an eigenvalue and both positive otherwise. When it is not necessary to distinguish between the two variants, we drop the superscript and just write $\operatorname{sep}_{\lambda}(A, B)$. For convenience, we also assume that $m \leq n$ throughout the paper.

^{*}Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, 39106 Germany. Current address: Department of Computer Science, Queens College / CUNY, 65-30 Kissena Boulevard, Flushing, NY 11367, USA, tim.mitchell@qc.cuny.edu. ORCID: 000-0002-8426-0242.

The two sep_{λ}(A, B) quantities can also be equivalently defined in terms of singular values as well as pseudospectra [GO06, pp. 348–349], where for some $\varepsilon \ge 0$, the ε -pseudospectrum of a matrix A is defined

$$\Lambda_{\varepsilon}(A) \coloneqq \{ z \in \mathbb{C} : z \in \Lambda(A + \Delta), \, \|\Delta\| \le \varepsilon \}, \tag{1.2a}$$

$$= \{ z \in \mathbb{C} : \sigma_{\min}(A - zI) \le \varepsilon \}.$$
(1.2b)

The first definition of pseudospectra dates to at least 1967, in Varah's Ph.D. thesis [Var67] with his introduction of an *r-approximate eigenvalue*, while in his 1979 paper on $\operatorname{sep}_{\lambda}^{V}(A, B)$, Varah used the term ε -spectrum for $\Lambda_{\varepsilon}(A)$. The current definitive reference on pseudospectra and their applications is certainly Trefethen and Embree's well-known book on the topic [TE05]. The term "pseudospectrum" was actually coined by Trefethen in 1990 [TE05, Ch. 6], 23 years after Varah's thesis, although it now considered the standard name.

The singular-value-based definitions of $sep_{\lambda}(A, B)$ are

$$\operatorname{sep}_{\lambda}^{\mathcal{V}}(A,B) = \min_{z \in \mathbb{C}} \{\sigma_{\min}(A - zI) + \sigma_{\min}(B - zI)\} = :\min_{z \in \mathbb{C}} f^{\mathcal{V}}(z),$$
(1.3a)

$$\operatorname{sep}_{\lambda}^{\mathrm{D}}(A,B) = \min_{z \in \mathbb{C}} \max\{\sigma_{\min}(A - zI), \, \sigma_{\min}(B - zI)\} \eqqcolon \min_{z \in \mathbb{C}} f^{\mathrm{D}}(z).$$
(1.3b)

For equivalent pseudospectral-based definitions of $sep_{\lambda}(A, B)$, we have

$$\operatorname{sep}_{\lambda}^{\mathcal{V}}(A,B) = \min_{\varepsilon_1, \varepsilon_2 \ge 0} \{ \varepsilon_1 + \varepsilon_2 : \Lambda_{\varepsilon_1}(A) \cap \Lambda_{\varepsilon_2}(B) \neq \emptyset \},$$
(1.4a)

$$\operatorname{sep}_{\lambda}^{D}(A,B) = \min_{\varepsilon \ge 0} \{ \varepsilon : \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B) \neq \emptyset \}.$$
(1.4b)

If $\varepsilon \geq \operatorname{sep}_{\lambda}^{D}(A, B)$ holds, then $\operatorname{int} \Lambda_{\varepsilon}(A) \cap \operatorname{int} \Lambda_{\varepsilon}(B) = \emptyset$ is a sufficient condition for $\varepsilon = \operatorname{sep}_{\lambda}^{D}(A, B)$. In contrast, while $\operatorname{int} \Lambda_{\varepsilon_{1}}(A) \cap \operatorname{int} \Lambda_{\varepsilon_{2}}(B) = \emptyset$ is a necessary condition for $\varepsilon_{1} + \varepsilon_{2} = \operatorname{sep}_{\lambda}^{V}(A, B)$ to hold, it is not a sufficient condition. This is because one can continuously adjust ε_{1} and ε_{2} such that the two pseudospectra always touch but never have interior points in common. For example, suppose that $\operatorname{sep}_{\lambda}^{V}(A, B) > 0$, and let $\hat{\varepsilon}_{1}$ be such that $\Lambda_{\hat{\varepsilon}_{1}}(A)$ and $\Lambda_{0}(B)$ only touch, i.e., an eigenvalue of B is in $\operatorname{bd} \Lambda_{\hat{\varepsilon}_{1}}(A)$ but $\operatorname{int} \Lambda_{\hat{\varepsilon}_{1}}(A) \cap \Lambda_{0}(B) = \emptyset$. In the same fashion, let $\hat{\varepsilon}_{2}$ be such that $\Lambda_{0}(A)$ and $\Lambda_{\hat{\varepsilon}_{2}}(B)$ only touch. Then by continuity of pseudospectra, it is clear that the 2D point $(\varepsilon_{1}, \varepsilon_{2})$ can be continuously adjusted between point $(\hat{\varepsilon}_{1}, 0)$ and point $(0, \hat{\varepsilon}_{2})$ such that $\operatorname{bd} \Lambda_{\varepsilon_{1}}(A) \cap \operatorname{bd} \Lambda_{\varepsilon_{2}}(B) \neq \emptyset$ and $\operatorname{int} \Lambda_{\varepsilon_{1}}(A) \cap \operatorname{int} \Lambda_{\varepsilon_{2}}(B) = \emptyset$ both always hold.

Varah called $\operatorname{sep}_{\lambda}^{V}(A, B)$ the spectrum separation in [Var79, Definition 3.2] due to its pseudospectral underpinnings; in fact, in his definition, he used the form given in (1.4a), not the other two alternatives. His motivation in defining $\operatorname{sep}_{\lambda}^{V}(A, B)$ was its connection to the sensitivity of solving the Sylvester equation:

$$AX - XB = C, (1.5)$$

where $X, C \in \mathbb{C}^{m \times n}$ and (1.5) has a unique solution if and only if A and B have no common eigenvalue. As Varah noted [Var79, p. 216], the sensitivity of a solution to (1.5) is inversely proportional to the separation of A and B:

$$\operatorname{sep}(A,B) \coloneqq \min_{\|X\|_{\mathrm{F}}=1} \|AX - XB\|_{\mathrm{F}} = \sigma_{\min}(I_n \otimes A - B^{\mathsf{T}} \otimes I_m),$$

a quantity which Stewart had earlier introduced for studying invariant subspaces [Ste73, Definition 4.5]. It holds that $0 \leq \operatorname{sep}(A, B) \leq \min_{\lambda \in \Lambda(A), \mu \in \Lambda(B)} |\lambda - \mu|$, and clearly, the lower bound is attained if and only if A and B have an eigenvalue in common, while the upper bound is attained if A and B are both normal. However, Varah stressed that if A or B is nonnormal, then $\operatorname{sep}(A, B)$ can be very close to zero, e.g., machine precision, even if the eigenvalues of A and B are well separated, and that $\operatorname{sep}(A, B)$ is often orders of magnitude smaller than $\operatorname{sep}_{\lambda}^{V}(A, B)$. In 1993, Higham's thorough error analysis for solving (1.5) numerically showed that bounding the error of a computed solution in terms of $\operatorname{sep}(A, B)^{-1}$ can sometimes "severely overestimate the effect of a perturbation on the data when only A and B are perturbed, because it does not take account of the special structure of the problem" [Hig93, p. 133], while simultaneously presenting an alternative error bound that remedies this deficiency. A few years later, Simoncini used $\operatorname{sep}_{\lambda}^{V}(A, B)$ and pseudospectra in her analysis of solving (1.5) via a Galerkin method [Sim96]. Meanwhile, Demmel initial interest in (his version of) $\operatorname{sep}_{\lambda}(A, B)$ was for problem of computing stable eigendecompositions [Dem83, Dem86], but in an entirely different context [Dem87], he subsequently used $\operatorname{sep}_{\lambda}^{D}(A, B)$ to disprove two conjectures respectively made by himself and Van Loan related to the (then unsolved) problem of computing the distance to instability of a stable matrix. Following in the spirit of using $\operatorname{sep}_{\lambda}(A, B)$ in the analysis of the stability of invariant subspaces of matrices [Var79, Dem83, Dem86], Karow and Kressner used $\operatorname{sep}_{\lambda}^{D}(A, B)$ in 2014 as a tool in deriving improved perturbation bounds [KK14]. Most recently in 2021, Roy et al. [RKBA21] used $\operatorname{sep}_{\lambda}^{V}(A, B)$ in connection with approximating pseudospectra of block triangular matrices; in this case, the value of $\operatorname{sep}_{\lambda}^{V}(A, B)$ can be used to construct several different outer approximations to pseudospectra of these structured matrices.

In terms of computing $\operatorname{sep}_{\lambda}(A, B)$, to the best of our knowledge, only a single algorithm has been given so far for $\operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$, due to Gu and Overton in 2006 [GO06], while no algorithms have appeared to date for $\operatorname{sep}_{\lambda}^{\mathrm{V}}(A, B)$. Nevertheless, computing $\operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ can at least approximate $\operatorname{sep}_{\lambda}^{\mathrm{V}}(A, B)$ to within a factor of two since

$$\frac{1}{2}\operatorname{sep}_{\lambda}^{\mathcal{V}}(A,B) \le \operatorname{sep}_{\lambda}^{\mathcal{D}}(A,B) \le \operatorname{sep}_{\lambda}^{\mathcal{V}}(A,B),$$
(1.6)

which is simply a special case of the relation $\frac{1}{n} \|x\|_1 \leq \|x\|_\infty \leq \|x\|_1$ for $x \in \mathbb{C}^n$ obtained by respectively identifying $\operatorname{sep}_{\lambda}^{\mathrm{V}}(A, B)$ and $\operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ with the 1-norm and ∞ -norm of $(\|E\|, \|F\|)^{\mathsf{T}} \in \mathbb{R}^2$. It is easy to obtain upper bounds for $\operatorname{sep}_{\lambda}(A, B)$ by simply evaluating f^{V} and/or f^{D} defined

in (1.3) at any points $z \in \mathbb{C}$, or better, by applying (nonsmooth) optimization techniques to find local minimizers of them. Due to the max function in $f^{\rm D}$, it is typically nonsmooth at minimizers, while f^{V} will be nonsmooth at a minimizer if that minimizer happens to coincide with an eigenvalue of A or B, which as Gu and Overton mentioned, is often the case for $\operatorname{sep}_{\lambda}^{V}(A, B)$. Despite the potential nonsmoothness, f^{V} and f^{D} are rather straightforward functions in just two real variables (via z = x + iy), whose function values and gradients (assuming z is a point where they are differentiable) can be obtained via computing $\sigma_{\min}(A-zI)$ and $\sigma_{\min}(B-zI)$ and their corresponding left and right singular vectors. When A and B are large and sparse, it is often still possible to efficiently compute f^{V} and f^{D} and their gradients via sparse methods. Nevertheless, finding local minimizers of (1.3) provides no guarantees for computing $\operatorname{sep}_{\lambda}(A, B)$, particularly since these problems may have many different local minima and the locally optimal function values associated with these minima may be very different. Moreover, in applications that use distances measures such as $sep_{\lambda}(A, B)$, obtaining an upper bound via local optimization is generally much less useful than either computing the actual measure or a lower bound to it. Indeed, in motivating their algorithm for $\operatorname{sep}_{\lambda}^{D}(A, B)$, Gu and Overton apply remarked [GO06, p. 350]: "the inability to verify global optimality [of minimizers of $f^{\rm D}$] remains a stumbling block preventing the computation of $sep_{\lambda}(A, B)$, or even the assessment of the quality of upper bounds, via optimization" and "in applications, lower bounds for such distance functions are more important than upper bounds, as they provide 'safety margins."

In this paper, we propose a new and much faster method to compute $\operatorname{sep}_{\lambda}^{D}(A, B)$ to arbitrary accuracy, using properties of pseudospectra, local optimization techniques, and a new methodology that we recently introduced in [Mit21] for finding global optimizers of singular value functions in two real variables. This new approach, called *interpolation-based globality certificates*, can be orders of magnitude faster than existing techniques and also avoids numerical difficulties inherent in older approaches. A modified version of our new $\operatorname{sep}_{\lambda}^{D}(A, B)$ algorithm also produces estimates of $\operatorname{sep}_{\lambda}^{V}(A, B)$ with stronger guarantees than those obtained by optimization; specifically, this modified method produces locally optimal upper bounds $\tilde{\varepsilon} = \varepsilon_1 + \varepsilon_2 \ge \operatorname{sep}_{\lambda}^{V}(A, B)$ such that $\operatorname{int} \Lambda_{\varepsilon_1}(A) \cap \operatorname{int} \Lambda_{\varepsilon_2}(B) = \emptyset$, which is a necessary condition for $\tilde{\varepsilon} = \operatorname{sep}_{\lambda}^{V}(A, B)$ to hold, but which optimization alone does not guarantee. Finally, we also propose a separate algorithm that is the first to compute $\operatorname{sep}_{\lambda}^{V}(A, B)$.

The paper is organized as follows. In §2, we give a brief overview of Gu and Overton's method for $\operatorname{sep}_{\lambda}^{D}(A, B)$ [GO06] and explain its shortcomings. Then, in §3, we give a high-level description of our new *optimization-with-restarts* method and an introduction to the ideas underlying interpolation-based globality certificates. As our new globality certificate for $\operatorname{sep}_{\lambda}^{D}(A, B)$ is quite different and significantly more complicated than those we devised for computing Kreiss constants and the distance to uncontrollability in [Mit21], we develop the necessary theoretical statements and components over three separate stages in §4, §5, and §6. In §7, we describe how to implement our completed algorithm and give its overall work complexity. We then turn to Varah's sep-lambda in §8. Numerical experiments are presented in §9, with concluding remarks given in §10.

2 Gu and Overton's method to compute $sep_{\lambda}^{D}(A, B)$ and its limitations

The algorithm of Gu and Overton for computing $\operatorname{sep}_{\lambda}^{D}(A, B)$ is particularly expensive: it is $\mathcal{O}((m+n)m^3n^3)$ work, e.g., $\mathcal{O}(n^7)$ when m = n, which makes it intractable for all but the tiniest of problems. The core of their method is a pair of related tests, each of which is inspired by a novel but expensive 2D level-set test developed earlier by Gu for estimating the distance to uncontrollability [Gu00]. The cost of each test is dominated by solving an associated generalized eigenvalue problem of order 4mn, which is $\mathcal{O}(m^3n^3)$ work when using standard dense eigensolvers.¹ Given some $\varepsilon \geq 0$, the first test ([GO06, Algorithm 1]) checks whether the ε -level sets of $\sigma_{\min}(A-zI)$ and $\sigma_{\min}(B-zI)$ have any points in common. If this is indeed the case, then clearly $\varepsilon \geq \operatorname{sep}_{\lambda}^{D}(A, B)$ must hold. However, if there are no level-set points in common, one cannot conclude that $\varepsilon < \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ holds. For example, having no shared level-set points may just be a consequence of $\Lambda_{\varepsilon}(A)$ being a subset of $\operatorname{int} \Lambda_{\varepsilon}(B)$ or vice versa, in which case, clearly $\varepsilon > \operatorname{sep}_{\lambda}^{D}(A, B)$ holds. To get around this difficulty, Gu and Overton devised an initialization procedure ([GO06, Algorithm 2]), which invokes their second test many times in order to compute an upper bound ε_{ub} such that for all $\varepsilon < \varepsilon_{ub}$, no connected component of $\Lambda_{\varepsilon}(A)$ can be strictly inside a component of $\Lambda_{\varepsilon}(B)$ or vice versa.² With this possibility excluded, i.e., $\varepsilon < \varepsilon_{\rm ub}$, the outcome of the first test then does indicates whether or not $\varepsilon < \sup_{\lambda}^{D}(A, B)$ holds. Gu and Overton's overall method [GO06, Algorithm 3] thus first computes ε_{ub} via their initialization procedure and then uses their first test to power a bisection iteration that converges to $\operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$. The entire bisection phase of their algorithm remains $\mathcal{O}(m^3n^3)$ work, since the number of bisection steps can be taken as a constant, but the initialization phase to compute the necessary ε_{ub} involves invoking the second test for (m+n) different parameter values, i.e., it solves (m+n) different generalized eigenvalue problems of order 4mn. Hence, the cost of their entire method is dominated by the initialization procedure, and the total asymptotic work complexity is $\mathcal{O}((m+n)m^3n^3)$.

In their concluding remarks [GO06, p. 358], Gu and Overton noted that the faster divide-andconquer technique of [GMO⁺06] for computing the distance to controllability could potentially be adapted to $\operatorname{sep}_{\lambda}^{D}(A, B)$, writing that "Although there are some inevitable difficulties with the numerical stability of this approach, the complexity drops significantly." Indeed, when m = n, adapting this divide-and-conquer approach would bring down the $\mathcal{O}(n^7)$ work complexity of their algorithm for $\operatorname{sep}_{\lambda}^{D}(A, B)$ to $\mathcal{O}(n^5)$ on average and $\mathcal{O}(n^6)$ in the worst case. However, this has not been implemented, and in our own experience of adapting this divide-and-conquer technique to other algorithms, we have observed that doing so can indeed come at the cost of significantly worse reliability due to numerical issues; see [Mit20, section 8].

Even with dense eigensolvers, Gu and Overton's method can be susceptible to numerical difficulties. A primary concern is that the first test (used for bisection) actually requires being able to assert whether or not two matrices have an eigenvalue in common. If eigenvalues can be computed exactly (which is possible in some cases, e.g., a diagonal matrix), then testing whether two matrices share an eigenvalue can be done without issues. However, in a practical code, computed eigenvalues will have rounding errors, and so one must generally resort to using a tolerance in order to carry out this test. But this also means that it is possible for the test to incorrectly assert that two eigenvalues are the same when they should only be considered close or vice versa. This is critical because the *binary decision* of bisection hinges upon the outcome of this numerical test.

 $^{^{1}}$ With respect to the usual convention of treating the computation of eigenvalues as an atomic operation with cubic work complexity, which we use throughout this paper.

²In [GO06], Gu and Overton state that this "not strictly inside" property holds for $\varepsilon \leq \varepsilon_{\rm ub}$, but actually this inequality should be strict. Near the top of [GO06, p. 354], it is claimed that " $\varepsilon = \sigma_{\min}(A - zI) > \sigma_{\min}(B - zI)$ " holds, where $z \in \operatorname{bd} \Lambda_{\varepsilon}(A)$ and $z \in \operatorname{int} \Lambda_{\varepsilon}(B)$. However, per [AGV17, p. 31], there can exist a finite number of points $z \in \operatorname{int} \Lambda_{\varepsilon}(B)$ such that $\sigma_{\min}(B - zI) = \varepsilon$, and so the "not strictly inside" claim may or may not hold when $\varepsilon = \varepsilon_{\rm ub}$. Fortunately, with inexact arithmetic, there is essentially no practical consequence of this small oversight, while the theory in [GO06] is corrected merely by replacing $\leq \operatorname{with} <$.

Making the wrong choice about the eigenvalues can cause bisection to erroneously update a lower or upper bound, which in turn can result in a significant or even complete loss of accuracy in the computed estimate. The distance-to-uncontrollability methods of [Gu00, BLO04, GMO⁺06] also have the same numerical pitfall. In the context of computing Kreiss constants via 2D level-set tests [Mit20], we recently proposed an improved procedure that does not require checking for shared eigenvalues, and as such, it is much more reliable in practice; see [Mit20, Key Remark 6.3]. Our improved technique can also be used to improve the reliability of the aforementioned distanceto-uncontrollability algorithms, but it does not appear to be applicable for Gu and Overton's algorithm for $\operatorname{sep}_{\lambda}^{D}(A, B)$. The fundamental difference in the $\operatorname{sep}_{\lambda}^{D}(A, B)$ setting is that Gu and Overton's first test is based upon checking whether or not the ε -level sets of two different functions, $\sigma_{\min}(A-zI)$ and $\sigma_{\min}(B-zI)$, have any points in common, whereas for the other quantities, pairs of points on a given level set of a single function are sought.

Finally, another way to provide some speedup to Gu and Overton's method would be to replace the bisection phase with an optimization-with-restarts iteration. In this case, a minimizer \tilde{z} of $f^{\rm D}$ with $\varepsilon = f^{\rm D}(\tilde{z})$ would be found using some nonsmooth optimization solver, and then, assuming $\varepsilon < \varepsilon_{ub}$, Gu and Overton's first test ([GO06, Algorithm 1]) would be used to assert whether or not \tilde{z} is a global minimizer of f^{D} . If so, then $\operatorname{sep}_{\lambda}^{D}(A, B) = \varepsilon$ and the computation is done. Otherwise, recalling that the first test computes the points $z \in \mathbb{C}$ such that $\sigma_{\min}(A-zI) = \sigma_{\min}(B-zI) = \varepsilon$, local optimization can be restarted from these points in order to find a better (lower) minimizer. Any such optimization-with-restarts method must monotonically converge to $\operatorname{sep}_{\lambda}^{D}(A, B)$ within a finite number of restarts because f^{D} only has a finite number of locally minimal function values, due to f^{D} being semialgebraic. However, there are some issues with this modification. The main limitation is that it neither accelerates nor removes the need for the initialization procedure for obtaining $\varepsilon_{\rm ub}$, which is $\mathcal{O}(n^7)$ work, while the subsequent convergent phase of either bisection or optimization-with-restarts is $\mathcal{O}(n^6)$ work. Consequently, any speedups will be both quite small and limited to the smallest values of n, while being essentially nonexistent for larger n. Another problem is that theory for nonsmooth optimization typically requires that solvers are initialized at points where the function is differentiable (see, e.g., [BCL⁺20, LO13, CMO17]), but by its nature, the points computed by Gu and Overton's first test are all points where $f^{\rm D}$ will almost certainly be nonsmooth. Hence, there may be issues in restarting optimization via these points, and depending on the exact solver and problem, we have observed that solvers can indeed stagnate at these initial points. It is not entirely clear how to best overcome this latter issue, but for us, it is not a priority. Instead, the focus of this paper is to propose an entirely different approach to computing $\operatorname{sep}^{D}_{\lambda}(A, B)$ that allows us to use optimization-with-restarts without any of the aforementioned drawbacks of extremely high costs, expensive initialization procedures, and various numerical and technical issues.

3 A high-level overview of our new $sep_{\lambda}^{D}(A, B)$ algorithm

To find a global minimizer of $f^{\rm D}$, a global optimization problem in two real variables, we will instead develop our optimization-with-restarts algorithm using *interpolation-based globality certificates* [Mit21]. The core task in developing such a method is to devise a generally continuous function (i.e., it may have some jumps) in one real variable that, given an estimate greater than the globally minimal value, has an identifiable subset of its domain with *positive measure* that provides a guaranteed way of locating new starting points for another round of optimization. When an estimate is globally minimal, this function should alternatively assert this fact somehow, e.g., by determining that the aforementioned subset is either empty or has measure zero. By sufficiently well approximating this function globally via a piecewise polynomial interpolant (this interpolant may also have jumps), e.g., by using Chebfun³ [DHT14], it is then possible to quickly check for the existence of the aforementioned positive measure subset, whose presence indicates that the estimate is not globally optimal. In fact, Chebfun can efficiently compute the precise set of intervals corresponding to this subset. When the estimate is too large, the property that there exists a subset of positive measure associated with new starting points is crucial for two reasons. First, it means that encountering this subset during the interpolation process is not a

³Available at https://www.chebfun.org.

probability zero event, and so if the function is well approximated, this subset will be detected. Second, optimization can be immediately restarted once any points in this subset are discovered, and so high-fidelity interpolants will often not be needed. As a result, restarts tend to be very inexpensive, while high-fidelity approximation is generally only needed for the final interpolant, which asserts that global convergence has indeed been obtained. Moreover, in practice only a handful of restarts are typically needed. Besides overall efficiency, interpolation-based globality certificates are inherently amenable to additional acceleration via parallel processing (see [Mit21, section 5.2]), while also being quite numerically robust compared to other techniques. There are several reasons for this latter property, but one is that by the nature of interpolation, global convergence is assessed as the result over many computations, whereas other approaches often rely upon a single computation that may result in an erroneous conclusion due to rounding errors; for more details, see [Mit21, sections 1.3 and 2.3].

Given some estimate $\varepsilon \geq \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$, in the next sections, we consider the problem of what function $d_{\varepsilon}: \mathbb{R} \to \mathbb{R}$ to devise for our globality certificate for either finding new points for restarting optimization or asserting whether $\varepsilon = \sup_{\lambda}^{D}(A, B)$ holds. The function d_{ε} should be reasonably well behaved and relatively cheap to evaluate, as otherwise approximating it could be prohibitively expensive and/or difficult. But as mentioned above, Chebfun can efficiently handle nonsmooth and discontinuous functions; [PPT09, pp. 905–906] describes the algorithm that Chebfun uses to efficiently detect discontinuities, either jumps in the function values or derivatives, which allows Chebfun to work around these difficult points during its approximation process. Consequently, we do not have to limit ourselves to smooth continuous candidates for d_{ε} . The function that we will propose is based on detecting whether or not int $\Lambda_{\varepsilon}(A) \cap \operatorname{int} \Lambda_{\varepsilon}(B)$ is empty and asserts that $\varepsilon > \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ if and only if $\min_{\theta \in (-\pi, \pi]} d_{\varepsilon}(\theta) < 0$ holds. Moreover, our certificate for detecting whether $\operatorname{int} \Lambda_{\varepsilon}(A) \cap \operatorname{int} \Lambda_{\varepsilon}(B) = \emptyset$ holds works for any value $\varepsilon > \operatorname{sep}_{\lambda}^{D}(A, B)$. As we will explain, our d_{ε} -based globality certificate incurs $\mathcal{O}(kn^3)$ work (recall that we assume $m \leq n$), where k is the number of function evaluations required to sufficiently approximate d_{ε} . Furthermore, our certificate also becomes more efficient the larger ε is, i.e., relatively few function evaluations are needed to approximate d_{ε} when $\varepsilon \gg \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ as compared to when $\varepsilon \approx \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$.

Remark 3.1. Although the first test of Gu and Overton ([GO06, Algorithm 1]) also detects if int $\Lambda_{\varepsilon}(A) \cap \operatorname{int} \Lambda_{\varepsilon}(B) = \emptyset$ holds, note that their test is both more limited in scope and more expensive than our d_{ε} -based certificate. Gu and Overton's first test (a) requires that $\varepsilon < \varepsilon_{ub}$ holds in order to use it, with ε_{ub} being very expensive to obtain, and (b) does the same amount of work regardless of the value of ε ; again, when m = n, computing ε_{ub} is $\mathcal{O}(n^7)$ work, while [GO06, Algorithm 1] is $\mathcal{O}(n^6)$ work.

4 Locating pseudospectral components

We now work on defining d_{ε} and establishing its properties, which is done over three sections. This section follows similarly to [Mit21, sections 2–4], where we first proposed interpolation-based globality certificates to find level-set components as tools for computing Kreiss constants and the distance to uncontrollability. Here we adapt these ideas to locating pseudospectral components, and throughout this section, we provide specific references to counterparts in [Mit21]. However, as will be seen, computing $\operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ is more complicated than computing these other quantities, and so the additional tools that we develop in §5 and §6 will also be needed.

Given a matrix $A \in \mathbb{C}^{m \times m}$, $\varepsilon \geq 0$, and some $z_0 \in \mathbb{C}$ such that ε is not a singular value of $A - z_0 I$, in this section we propose a way of determining which rays emanating from z_0 intersect with $\Lambda_{\varepsilon}(A)$ and which do not. We define the ray emanating from z_0 specified by angle $\theta \in \mathbb{R}$ as

$$\mathcal{R}_{\theta} \coloneqq \{ z_0 + r \mathrm{e}^{\mathrm{i}\theta} \in \mathbb{C} : r \in \mathbb{R}, \, r > 0 \}.$$

$$(4.1)$$

As we will explain momentarily, our assumption on ε ensures that a condition needed by our method indeed holds; relatedly, our assumption also ensures that the "search point" z_0 is not on the boundary of $\Lambda_{\varepsilon}(A)$. Consider the following function parameterized in polar coordinates:

$$f_A(r,\theta) = \sigma_{\min}(A - (z_0 + re^{\mathbf{i}\theta})I)) = \sigma_{\min}(F_A(r,\theta)), \text{ where } F_A(r,\theta) = \mathbf{i}e^{-\mathbf{i}\theta}(A - z_0I) - \mathbf{i}rI \quad (4.2)$$

and the second equality above holds since multiplication by a unitary scalar does not alter the singular values of a matrix. Note that $\mathrm{bd} \Lambda_{\varepsilon}(A)$ is contained in the ε -level set of f_A . The following pair of results give us a way to determine whether or not \mathcal{R}_{θ} and $\Lambda_{\varepsilon}(A)$ intersect, and when they do, to also calculate all the points in $\mathcal{R}_{\theta} \cap \mathrm{bd} \Lambda_{\varepsilon}(A)$. The first of these two results is yet another variation of the 1D level-set technique Byers introduced in order to develop the first method for computing the distance to instability in 1988 [Bye88], a powerful tool which we and many others have adapted, extended, or used to develop 1D and 2D methods to compute various quantities. Applications include the \mathcal{H}_{∞} and \mathcal{L}_{∞} norms [BBK89, BB90, BS90, BSV12, BM18], distance to uncontrollability [Bye90, GN93, Gu00, Mit21], numerical radius [HW97, MO05, Mit22], pseudospectral (or spectral value set) abscissa and radius [BLO03, MO05, BM19], Kreiss constants [Mit20, Mit21], as well as the optimization of passive systems [MVD20b, MVD20a, MVD20a].

Lemma 4.1 (cf. [Mit21, Theorems 2.1, 3.1, and 4.1]). Let $A \in \mathbb{C}^{m \times m}$, $\varepsilon \geq 0$, $z_0 \in \mathbb{C}$, and $r, \theta \in \mathbb{R}$. Then ε is a singular value of $F_A(r, \theta)$ defined in (4.2) if and only if ir is an eigenvalue of the Hamiltonian matrix

$$C_{\theta} \coloneqq \begin{bmatrix} \mathbf{i} \mathrm{e}^{-\mathbf{i}\theta} (A - z_0 I) & -\varepsilon I \\ \varepsilon I & \mathbf{i} \mathrm{e}^{\mathbf{i}\theta} (A - z_0 I)^* \end{bmatrix}.$$
(4.3)

Proof. Suppose that ε is a singular value of $F_A(r, \theta)$ with left and right singular vectors u and v. Then

$$\varepsilon \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} F_A(r,\theta) & 0 \\ 0 & F_A(r,\theta)^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = \begin{bmatrix} \mathbf{i}e^{-\mathbf{i}\theta}(A-z_0I) & 0 \\ 0 & -\mathbf{i}e^{\mathbf{i}\theta}(A-z_0I)^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} + \mathbf{i}r \begin{bmatrix} -I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix}.$$

Rearranging terms, using the fact that $\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix}$, and multiplying the bottom block row by -1, we obtain $C_{\theta}\begin{bmatrix} v \\ u \end{bmatrix} = \mathbf{i}r\begin{bmatrix} v \\ u \end{bmatrix}$.

Corollary 4.2. Let $A \in \mathbb{C}^{m \times m}$, $\varepsilon \geq 0$, $z_0 \in \mathbb{C}$, $r, \theta \in \mathbb{R}$, and \mathcal{R}_{θ} be the ray defined by (4.1). Then $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset$ if and only if **i**r is an eigenvalue of C_{θ} with r > 0.

Proof. Suppose that \mathcal{R}_{θ} and $\Lambda_{\varepsilon}(A)$ intersect. As $\Lambda_{\varepsilon}(A)$ is bounded, there exists an r > 0 such that the point $z_0 + re^{i\theta}$ is also on the boundary of $\Lambda_{\varepsilon}(A)$, and so $\sigma_{\min}(F_A(r,\theta)) = \varepsilon$. Thus by Lemma 4.1, ir is an eigenvalue of C_{θ} . Now suppose C_{θ} has some eigenvalue ir with r > 0. Again by Lemma 4.1, ε must then be a singular value of $F_A(r,\theta)$ but not necessarily the smallest one. Thus, $\sigma_{\min}(F_A(r,\theta)) = \hat{\varepsilon} \leq \varepsilon$ and so it follows that $z_0 + re^{i\theta}$ is in $\Lambda_{\hat{\varepsilon}}(A) \subseteq \Lambda_{\varepsilon}(A)$.

For any $z_0 + re^{i\theta} \in bd \Lambda_{\varepsilon}(A)$ with r > 0, clearly $\sigma_{\min}(F_A(r,\theta)) = \varepsilon$, and so by Lemma 4.1, i $r \in \Lambda(C_{\theta})$. Hence, via computing all of the imaginary eigenvalues of C_{θ} , Lemma 4.1 provides a way to calculate all of the points in $\mathcal{R}_{\theta} \cap bd \Lambda_{\varepsilon}(A)$. However, note that if $ir \in \Lambda(C_{\theta})$ with r > 0, then $z_0 + re^{i\theta}$ may or may not be on $bd \Lambda_{\varepsilon}(A)$. There are two reasons for this. First, per the proof of Corollary 4.2, ε may not be the smallest singular value of $F_A(r,\theta)$, in which case $z_0 + re^{i\theta} \in \Lambda_{\varepsilon}(A)$ for some $\hat{\varepsilon} < \varepsilon$. Second, there can exist a *finite number* of points $z \in \Lambda_{\varepsilon}(A)$ such that $z \notin bd \Lambda_{\varepsilon}(A)$ but $\sigma_{\min}(F_A(r,\theta)) = \varepsilon$ nevertheless holds; see [AGV17, p. 31].

Corollary 4.2 can be stated more strongly, i.e., in terms of a line intersecting $\Lambda_{\varepsilon}(A)$, since ir_{neg} with $r_{\text{neg}} < 0$ is an eigenvalue of C_{θ} if and only if $\mathbf{i}|r_{\text{neg}}|$ is an eigenvalue of $C_{\theta+\pi}$. However, for developing the theoretical concepts for our algorithm, it will be more intuitive and simpler to work with the notion of rays emanating from z_0 for the time being. For a code, it does make sense to take advantage of all the imaginary eigenvalues of C_{θ} , and we describe how this is done, along with other implementation details, in §7. Regarding the spectrum of C_{θ} , also note that since C_{θ} is Hamiltonian, its eigenvalues are symmetric with respect to the imaginary axis. Eigenvalues of *real* Hamiltonian matrices have additional symmetry with respect to the real axis, but this is generally not the case for the spectrum of C_{θ} due to C_{θ} being generically complex valued (even if A is real). Structure-preserving eigensolvers exist, e.g., [BMX99], that preserve this eigenvalue symmetry numerically. **Remark 4.3.** While Lemma 4.1 pertains to the eigenvalues of a single Hamiltonian matrix, the analogous [Mit21, Theorems 2.1, 3.1, and 4.1] used for computing Kreiss constants and the distance to uncontrollability via interpolation-based globality certificates are in terms of the eigenvalues of certain structured matrix pencils. For the case of Kreiss constants [Mit21, Theorems 2.1 and 3.1], the associated matrix pencils include parametric matrices that can be singular, and so these generalized eigenvalue problems cannot be reduced to standard eigenvalue problems. However, the matrix pencil for the distance to uncontrollability does permit such a reduction, i.e., [Mit21, Theorem 4.1] can be simplified to be in terms of the eigenvalues of the complex Hamiltonian matrix

$$\begin{bmatrix} i e^{-i\theta} A & \widetilde{B} \\ \gamma I & i e^{i\theta} A^* \end{bmatrix},$$
(4.4)

where in (4.4), matrices A and \tilde{B} and scalars γ and θ are defined in [Mit21, Theorem 4.1].

As we will soon see, we will need to preclude the possibility of zero being an eigenvalue of C_{θ} . The following straightforward result shows that our assumption on ε not being a singular value of $A - z_0 I$ accomplishes this.

Lemma 4.4 (cf. [Mit21, Theorems 2.4, 3.3, and 4.4]). Let $A \in \mathbb{C}^{m \times m}$, $\varepsilon \in \mathbb{R}$, $z_0 \in \mathbb{C}$, and $\theta \in \mathbb{R}$. Then the matrix C_{θ} defined in (4.3) has zero as an eigenvalue if and only if the matrix $(A - z_0 I)(A - z_0 I)^*$ has ε^2 as an eigenvalue.

Proof. Since the blocks of C_{θ} are all square matrices of the same size, and the lower two blocks, εI and $ie^{i\theta}(A - z_0 I)^*$, commute, we have that

$$\det(C_{\theta}) = \det(-(A - z_0 I)(A - z_0 I)^* - (-\varepsilon I)(\varepsilon I)) = \det((A - z_0 I)(A - z_0 I)^* - \varepsilon^2 I),$$

thus proving the if-and-only-if equivalence.

We are now ready to present the first major component in our construction of d_{ε} . Given $\varepsilon \ge 0$ specifying the ε -pseudospectrum of A, and $z_0 \in \mathbb{C}$ such that ε is not a singular value of $A - z_0 I$, we define the function $a_{\varepsilon} : (-\pi, \pi] \to [0, \pi^2]$ and associated set (cf. [Mit21, Equations (2.4), (3.4), and (4.4)]):

$$a_{\varepsilon}(\theta) \coloneqq \min\{\operatorname{Arg}(-\mathbf{i}\lambda)^2 : \lambda \in \Lambda(C_{\theta}), \operatorname{Re}\lambda \le 0\},\tag{4.5a}$$

$$\mathcal{A}_{\varepsilon} \coloneqq \{\theta : a_{\varepsilon}(\theta) = 0, \, \theta \in (-\pi, \pi]\},\tag{4.5b}$$

where $\operatorname{Arg} : \mathbb{C} \setminus \{0\} \to (-\pi, \pi]$ is the principal value argument function, the matrix C_{θ} is defined in (4.3), and the term $\operatorname{Arg}(-\mathbf{i}\lambda)$ in (4.5a) is squared in order to smooth its value out when transitioning to/from zero. We explain this in more detail later on, but the squaring is done in order to make a_{ε} easier to approximate globally on its domain. Note that the definition of a_{ε} excludes eigenvalues in the open right half of the complex plane since the spectrum of C_{θ} is symmetric with respect to the imaginary axis.

Theorem 4.5 (Properties of a_{ε} ; cf. [Mit21, Theorems 2.7, 3.5, and 4.6]). Let $A \in \mathbb{C}^{m \times m}$, $\varepsilon \geq 0$, and $z_0 \in \mathbb{C}$ be such that ε is not a singular value of $A - z_0 I$. Then, the function a_{ε} defined in (4.5a) has the following properties:

- (i) $a_{\varepsilon}(\theta) \geq 0$ on its entire domain, i.e., $\forall \theta \in (-\pi, \pi]$,
- (ii) $a_{\varepsilon}(\theta) = 0 \quad \iff \quad \exists r > 0 \text{ such that } \mathbf{i}r \in \Lambda(C_{\theta}) \quad \iff \quad \mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset,$
- (iii) a_{ε} is continuous on its entire domain,
- (iv) a_{ε} is differentiable at a point θ if the eigenvalue $\lambda \in \Lambda(C_{\theta})$ attaining the value of $a_{\varepsilon}(\theta)$ is unique and simple.

Furthermore, the following properties hold for the associated set $\mathcal{A}_{\varepsilon}$ defined in (4.5b):

(v)
$$\varepsilon = 0 \iff \mu(\mathcal{A}_{\varepsilon}) = 0$$
,

- (vi) $\varepsilon_1 < \varepsilon_2 \iff \mu(\mathcal{A}_{\varepsilon_1}) < \mu(\mathcal{A}_{\varepsilon_2}),$
- (vii) if $\varepsilon > f_A(0,\theta)$ for any $\theta \in \mathbb{R}$, then $\mu(\mathcal{A}_{\varepsilon}) = 2\pi$,
- (viii) $\mathcal{A}_{\varepsilon}$ can have up to m connected components.

Proof. Noting that $-\mathbf{i}\lambda$ in (4.5a) is always in the (closed) upper half of the complex plane, statements (i) and (ii) hold by the definition of a_{ε} and Corollary 4.2. Statement (iii) follows from the continuity of eigenvalues and our assumption that ε is not a singular value of $A - z_0 I$, equivalently $\varepsilon^2 \notin \Lambda((A - z_0 I)(A - z_0 I)^*)$, and thus, by Lemma 4.4, $0 \notin \Lambda(C_{\theta})$ is ensured for any θ . Statement (iv) follows from standard perturbation theory for simple eigenvalues and by the definition of a_{ε} .

Now turning to $\mathcal{A}_{\varepsilon}$, either $z_0 \in \operatorname{int} \Lambda_{\varepsilon}(A)$ or $z_0 \notin \Lambda_{\varepsilon}(A)$ must hold since our assumption on ε precludes z_0 from being a boundary point. If $\varepsilon > f_A(0,\theta)$, then $z_0 \in \operatorname{int} \Lambda_{\varepsilon}(A)$, which in turn implies that $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset$ for all θ , thus proving (vii). Now assume $z_0 \notin \Lambda_{\varepsilon}(A)$. Statement (viii) is a consequence of the well-known fact that for any matrix $A \in \mathbb{C}^{m \times m}$, its ε -pseudospectrum has at most m connected components. For any component \mathcal{G} of $\Lambda_{\varepsilon}(A)$, by connectedness and (ii), it is clear that \mathcal{G} is associated with a single interval $\mathcal{I} \subseteq (-\pi, \pi]$ such that $a_{\varepsilon}(\theta) = 0$ if and only if $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset$. Since $\mathcal{A}_{\varepsilon}$ is simply the union of those intervals associated with the components of $\Lambda_{\varepsilon}(A)$, of which there can be at most m, $\mathcal{A}_{\varepsilon}$ also has at most m components, thus proving (viii). Statement (vi) follows by noting that $\Lambda_{\varepsilon_1}(A) \subset \Lambda_{\varepsilon_2}(A)$ is equivalent to $\mathcal{A}_{\varepsilon_1} \subset \mathcal{A}_{\varepsilon_2}$. Since $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon_1}(A) \neq \emptyset$ implies that $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon_2}(A) \neq \emptyset$, it follows that $a_{\varepsilon_1}(\theta) = 0$ implies $a_{\varepsilon_2}(\theta) = 0$, and so $\mathcal{A}_{\varepsilon_1} \subset \mathcal{A}_{\varepsilon_2}$. Now suppose that $\mathcal{A}_{\varepsilon_1} \supset \mathcal{A}_{\varepsilon_2}$ and let $\theta \in \mathcal{A}_{\varepsilon_1} \setminus \mathcal{A}_{\varepsilon_2}(A)$ cannot hold, a contradiction. Finally, for (v), if $\varepsilon = 0$, $\Lambda_{\varepsilon}(A) = \Lambda(A)$, and so $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset$ for at most mdifferent angles. As there can be at most m connected components of $\mathcal{A}_{\varepsilon}$, if $\mu(\mathcal{A}_{\varepsilon}) = 0$ holds, then $\varepsilon = 0$.

Per Theorem 4.5, a_{ε} is a continuous function and $a_{\varepsilon}(\theta) = 0$ if and only if $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset$. Thus, by finding roots of a_{ε} , we find rays which intersect the ε -pseudospectrum of A, our first step toward finding regions where $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ overlap. For an illustration of this correspondence, see Fig. 4.1, where b_{ε} , the analogue of a_{ε} for matrix B, is also plotted.

The properties of a_{ε} listed in Theorem 4.5 show that it is reasonably well behaved. Satisfying the assumption that ε is not a singular value of $A-z_0I$ can be trivially met, e.g., just by choosing z_0 with a bit of randomness. The dominant cost of evaluating a_{ε} at a point θ is computing the spectrum of C_{θ} , i.e., $\mathcal{O}(m^3)$ work. Relative to Gu and Overton's $\operatorname{sep}_{\lambda}^{D}(A, B)$ algorithm, this is a negligible cost. To find the roots of a_{ε} , we can approximate a_{ε} using Chebfun, which is why we defined a_{ε} using the squared term $\operatorname{Arg}(-i\lambda)^2$ instead of just $\operatorname{Arg}(-i\lambda)$. As will be made clear in §6, a_{ε} transitioning to/from zero corresponds to two (or possibly more non-generically) eigenvalues of C_{θ} coalescing on the positive portion of the imaginary axis. Without this squaring, a_{ε} would generally be non-Lipschitz at such transition points and thus it could be difficult and/or expensive to approximate via interpolation; the squaring smooths out this high rate of change so that a_{ε} is easier to approximate. Although the analogues of a_{ε} and Theorem 4.5 that appeared in [Mit21] for computing Kreiss constants and the distance to uncontrollability were sufficient to develop interpolation-based globality certificates for those quantities, for $\operatorname{sep}_{\lambda}^{D}(A, B)$, a_{ε} and Theorem 4.5 are insufficient.

5 Locating pseudospectral overlap

As part of locating regions where $\Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B) \neq \emptyset$, we will also need to locate the components of $\Lambda_{\varepsilon}(B)$ with respect to the same "search point" z_0 and given value of ε . Thus, for matrix B, let f_B and F_B respectively denote the analogues of f_A and F_A defined in (4.2), and similarly, let b_{ε} and $\mathcal{B}_{\varepsilon}$ be respective analogues of a_{ε} and $\mathcal{A}_{\varepsilon}$ defined in (4.5). For matrix A, we continue to use C_{θ} to denote its associated Hamiltonian matrix defined in (4.3), while we use S_{θ} to denote the analogue Hamiltonian matrix for B, as both matrices will be needed. Per the assumption of Theorem 4.5, we now need to assume that ε is not a singular value of either $A - z_0I$ or $B - z_0I$,



Figure 4.1: For two randomly generated matrices $A, B \in \mathbb{C}^{14 \times 14}$, the left pane shows their eigenvalues (respectively x's and dots), and $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ (respectively solid and dotted contours) for $\varepsilon = 0.3 > \sup_{\lambda}^{D}(A, B)$. The search point z_{0} is the origin; rays emanating from it are depicted by dashed lines. The right pane shows corresponding plots of $a_{\varepsilon}, b_{\varepsilon}$ (respectively solid and dotted curves), and ℓ_{ε} (dashed), where a_{ε} is defined in (4.5a), b_{ε} is its analogue for matrix B, and ℓ_{ε} is defined in (5.1a). On the left, rays in the lower left quadrant only intersect $\Lambda_{\varepsilon}(A)$ or neither ε -pseudospectrum, while rays in the lower right quadrant only intersect $\Lambda_{\varepsilon}(B)$ or neither ε -pseudospectrum. Correspondingly, for $(-\pi, -\frac{1}{2}\pi]$ on the right, we see that a_{ε} has zeros but b_{ε} is always positive, and vice versa for $(-\frac{1}{2}\pi, 0]$. Meanwhile, there exist rays in the upper right quadrant that pass through both $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$, but $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ do not overlap in this region; thus, on the right for $(0, \frac{1}{2}\pi]$, we see that a_{ε} and b_{ε} do have zeros in common, but ℓ_{ε} is equal to zero on this interval. Finally, in the upper left quadrant, int $\Lambda_{\varepsilon}(A)$ and int $\Lambda_{\varepsilon}(B)$ do in fact overlap, and so on the right, we see that a_{ε} and b_{ε} have zeros in common and ℓ_{ε} is indeed negative on a subset of $(\frac{1}{2}\pi, \pi]$ with positive measure.

which again, can be easily satisfied by choosing z_0 with some randomness. In establishing tools for locating pseudospectral overlap, we will make use of the following elementary result.

Lemma 5.1. Let $\mathcal{A}, \mathcal{B} \subset \mathbb{R}$ be such that \mathcal{A} and \mathcal{B} respectively consist of m and n connected components. Then $\mathcal{A} \cap \mathcal{B}$ can have up to m + n - 1 connected components.

Proof. Let $\mathcal{A} = \mathcal{A}_1 \cup \cdots \cup \mathcal{A}_m$, where each \mathcal{A}_j is a connected component of \mathcal{A} and $\mathcal{A}_j \cap \mathcal{A}_k = \emptyset$ for all $j \neq k$, and in an analogous fashion, let $\mathcal{B} = \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_n$. Without loss of generality, assume that $m \leq n$. If m = 1, suppose that the claim is not true, i.e., that $\mathcal{A} \cap \mathcal{B}$ has more than n components. Then there exists at least one pair of numbers x and y that are in different components of $\mathcal{A} \cap \mathcal{B}$ but must be in the same component \mathcal{B}_i of \mathcal{B} . However, by connectedness of the components of \mathcal{A} and \mathcal{B} , we have that $[x, y] \subset \mathcal{A}_1 = \mathcal{A}$ and $[x, y] \subset \mathcal{B}_j$. Therefore $[x, y] \subset \mathcal{A} \cap \mathcal{B}$, contradicting that x and y are in different components of $\mathcal{A} \cap \mathcal{B}$. For the inductive step, now assume that the claim holds when \mathcal{A} consists of j components, for $j = 1, \ldots, m-1$ and j < n, and suppose that \mathcal{A} has m components. Let $s = \frac{1}{2}(a_{\rm L} + a_{\rm R})$, where $a_{\rm L} = \sup_{a \in \mathcal{A}_{m-1}} a$ and $a_{\rm R} = \inf_{a \in \mathcal{A}_m} a$, and define $\mathcal{B}_{L} := \{b : b \in \mathcal{B}, b < s\}$ and $\mathcal{B}_{R} := \{b : b \in \mathcal{B}, b > s\}$. Clearly \mathcal{B}_{L} and \mathcal{B}_{R} are disjoint and $\mathcal{B}_{L} \cup \mathcal{B}_{R} = \mathcal{B} \setminus \{s\}$. Letting n_{L} and n_{R} denote the respective number of connected components of \mathcal{B}_{L} and \mathcal{B}_{R} , it follows that $n_{L} + n_{R} = n$ if $s \notin int \mathcal{B}$ and $n_{L} + n_{R} = n + 1$ otherwise. Applying the inductive hypothesis, $\{A_1 \cup \cdots \cup A_{m-1}\} \cup B_L$ has at most $(m-1) + n_L - 1$ connected components, while $\mathcal{A}_m \cup \mathcal{B}_R$ has at most n_R connected components. Noting that $\{\mathcal{A}_1 \cup \cdots \cup \mathcal{A}_{m-1}\} \cup \mathcal{B}_L$ and $\mathcal{A}_m \cup \mathcal{B}_R$ are also disjoint and their union is $\mathcal{A} \cap \mathcal{B}$, since $s \notin \mathcal{A} \cap \mathcal{B}$, it follows that $\mathcal{A} \cap \mathcal{B}$ has at most $(m-1) + n_{\rm L} - 1 + n_{\rm R} \le m + n - 1$ connected components. The bound is tight, as one can construct \mathcal{A} such that \mathcal{A}_j intersects both \mathcal{B}_j and \mathcal{B}_{j+1} for $j = 1, \ldots, m-1$, while \mathcal{A}_m intersects \mathcal{B}_j for $j = m - 1, \ldots, n$. **Theorem 5.2** (Properties of $a_{\varepsilon} + b_{\varepsilon}$ and a necessary condition for overlap). Let $A \in \mathbb{C}^{m \times m}$, $B \in \mathbb{C}^{n \times n}$, $\varepsilon \geq 0$, and $z_0 \in \mathbb{C}$ be such that ε is not a singular value of either $A - z_0I$ or $B - z_0I$, and let \mathcal{R}_{θ} be the ray defined in (4.1). Furthermore, let $\mathcal{Z}_{\varepsilon} := \{\theta \in (-\pi, \pi] : a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = 0\}$, where a_{ε} is defined in (4.5a) for A and b_{ε} is its analogue for B. Then the following statements hold:

- (i) if $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B) \neq \emptyset$, then $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = 0$,
- (ii) if $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = 0$, then $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset$ and $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B) \neq \emptyset$,
- (iii) $a_{\varepsilon} + b_{\varepsilon}$ is continuous on its entire domain $(-\pi, \pi]$,
- (iv) $a_{\varepsilon} + b_{\varepsilon}$ is differentiable at a point θ if a_{ε} and b_{ε} are differentiable at θ ,
- (v) $\mathcal{Z}_{\varepsilon}$ can have up to m + n 1 connected components.

Proof. The assumption in (i) implies that $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \neq \emptyset$ and $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B) \neq \emptyset$ hold, and so $a_{\varepsilon}(\theta) = 0$ and $b_{\varepsilon}(\theta) = 0$ by Theorem 4.5 (ii). Statements (ii)–(iv) are direct consequences of Theorem 4.5 (ii)–(iv). For (v), note that $\mathcal{Z}_{\varepsilon} = \mathcal{A}_{\varepsilon} \cap \mathcal{B}_{\varepsilon}$, where $\mathcal{A}_{\varepsilon}$ is defined in (4.5b) for Aand $\mathcal{B}_{\varepsilon}$ is its analogue for B. As $\mathcal{A}_{\varepsilon}$ and $\mathcal{B}_{\varepsilon}$ respectively have up to m and n connected components by Theorem 4.5, statement (v) follows from Lemma 5.1.

Given an angle θ , Theorem 5.2 states that $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = 0$ is a necessary condition for the pseudospectra $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ to overlap somewhere along the ray \mathcal{R}_{θ} , but it is easy to see that this is not a sufficient condition for such overlap. To obtain such a sufficient condition, we now define the function $\ell_{\varepsilon} : (-\pi, \pi] \to (-\infty, 0]$ and an associated set:

$$\ell_{\varepsilon}(\theta) \coloneqq -\mu \left(\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B) \right), \qquad (5.1a)$$

$$\mathcal{L}_{\varepsilon} \coloneqq \{\theta \in (-\pi, \pi] : \ell_{\varepsilon}(\theta) < 0\}.$$
(5.1b)

As $\mathcal{L}_{\varepsilon}$ is open, it is measurable, and via Lemma 4.1, we know that the ray \mathcal{R}_{θ} can intersect at most 2m and 2n boundary points, respectively, of $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$. Thus, the number of connected components of $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A)$ is finite, as is the number of connected components of $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B)$; hence, the intersection in the definition of ℓ_{ε} is measurable. Moreover, Lemma 4.1 allows us to determine these intervals (or isolated points), and so the value of $\ell_{\varepsilon}(\theta)$ can be computed simply by calculating how much the intervals of $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A)$ overlap those of $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B)$; we explain exactly how this is done in §7. In addition to a_{ε} and b_{ε} , function ℓ_{ε} is also plotted in Fig. 4.1.

Theorem 5.3 (Properties of ℓ_{ε} and a sufficient condition for overlap). Let $A \in \mathbb{C}^{m \times m}$, $B \in \mathbb{C}^{n \times n}$, $\varepsilon \geq 0, z_0 \in \mathbb{C}, \theta \in \mathbb{R}$, and \mathcal{R}_{θ} be the ray defined in (4.1). Then for the function ℓ_{ε} defined in (5.1a), the following statements hold:

- (i) $\ell_{\varepsilon}(\theta) < 0 \iff \mathcal{R}_{\theta} \cap \operatorname{int} \Lambda_{\varepsilon}(A) \cap \operatorname{int} \Lambda_{\varepsilon}(B) \neq \emptyset$,
- (ii) if $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) > 0$, then $\ell_{\varepsilon}(\theta) = 0$,
- (iii) ℓ_{ε} is continuous on its entire domain $(-\pi, \pi]$,
- (iv) ℓ_{ε} is differentiable at a point θ if $\forall r > 0$ such that $z_0 + re^{i\theta} \in bd \Lambda_{\varepsilon}(A)$, ir is a simple eigenvalue of C_{θ} , and $\forall r > 0$ such that $z_0 + re^{i\theta} \in bd \Lambda_{\varepsilon}(B)$, ir is a simple eigenvalue of S_{θ} .

Furthermore, the following statements hold for the associated set $\mathcal{L}_{\varepsilon}$ defined in (5.1b):

- (v) $\varepsilon \leq \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B) \iff \mu(\mathcal{L}_{\varepsilon}) = 0,$
- (vi) $\operatorname{sep}_{\lambda}^{D}(A, B) < \varepsilon_{1} < \varepsilon_{2} \quad \Longleftrightarrow \quad 0 < \mu(\mathcal{L}_{\varepsilon_{1}}) < \mu(\mathcal{L}_{\varepsilon_{2}}),$
- (vii) $\min_{\theta \in (-\pi,\pi]} \ell_{\varepsilon}(\theta) < 0 \iff 0 < \mu(\mathcal{L}_{\varepsilon}) \iff \operatorname{sep}_{\lambda}^{\mathrm{D}}(A,B) < \varepsilon.$

Proof. Statement (i) simply follows from the definition of ℓ_{ε} given in (5.1a) and noting that the intersection $\mathcal{R}_{\theta} \cap \operatorname{int} \Lambda_{\varepsilon}(A) \cap \operatorname{int} \Lambda_{\varepsilon}(B)$ is either empty or consists of a finite number of open intervals in \mathbb{R} . For (ii), if $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) > 0$, then either $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) = \emptyset$ or $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B) = \emptyset$ holds by Theorem 4.5 (ii), and so $\ell_{\varepsilon}(\theta) = 0$. Statement (iii) follows from the fact the boundaries of ε -pseudospectra vary continuously with respect to ε , which is clear from (1.2), and via Lemma 4.1, do not contain any straight line segments. Under the assumptions in (iv), standard perturbation theory for simple eigenvalues applies.

For $\mathcal{L}_{\varepsilon}$, (v) is a direct consequence of (i) and the definition of $\operatorname{sep}_{\lambda}^{D}(A, B)$ given in (1.4b), as $\operatorname{int} \Lambda_{\varepsilon}(A) \cap \operatorname{int} \Lambda_{\varepsilon}(B) = \emptyset$ if and only if $\varepsilon \leq \operatorname{sep}_{\lambda}^{D}(A, B)$. Statement (vi) follows by a similar argument to the proof of Theorem 4.5 (vi), with $\mu(\mathcal{L}_{\varepsilon_{1}}) > 0$ if and only if $\varepsilon_{1} > \operatorname{sep}_{\lambda}^{D}(A, B)$ following from (i). Statement (vii) is simply a combination of (i) and (vi).

From Theorem 5.3 (vii), it is clear that if ℓ_{ε} can be sufficiently well approximated, then one can determine whether or not $\varepsilon > \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ holds. Moreover, as we fully explain in §7, via Lemma 4.1, knowledge of such angles can be used to compute points on the ε -level set of f^{D} , points which can be used to restart optimization to find a better (lower) estimate for $\operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$. Thus, one may wonder what the point was of considering $a_{\varepsilon} + b_{\varepsilon}$ and deriving its associated necessary condition given in Theorem 5.2. There is in fact a very important reason for this.

As ℓ_{ε} is constant (zero) whenever it is not negative, it can, ironically, be a difficult function to approximate. The pitfall here is that regions where a function appears to be constant may be *undersampled* by interpolation software, precisely because the computed estimate of the error on such regions will generally be exactly zero, e.g., because the software initially builds a constant interpolant for the region in question. Thus, there is a concern that approximating ℓ_{ε} via interpolation may miss regions where $\ell_{\varepsilon}(\theta) < 0$ holds, particularly if these regions are small compared to the regions where $\ell_{\varepsilon}(\theta) = 0$. Our solution to this difficulty is to replace ℓ_{ε} by another non-constant function whenever $\ell_{\varepsilon}(\theta) = 0$ holds. We first consider the continuous function $t_{\varepsilon} : (-\pi, \pi] \to \mathbb{R}$

$$t_{\varepsilon}(\theta) \coloneqq \begin{cases} a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) & \text{if } a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) > 0\\ \ell_{\varepsilon}(\theta) & \text{otherwise} \end{cases},$$
(5.2a)

$$\mathcal{T}_{\varepsilon} \coloneqq \{\theta \in (-\pi, \pi] : t_{\varepsilon}(\theta) = 0\},\tag{5.2b}$$

an alternative to approximating ℓ_{ε} ; we have also defined $\mathcal{T}_{\varepsilon}$, the set of roots of t_{ε} , as this will be used later. The key point here is that t_{ε} tells us at which angles the sufficient condition for $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ to overlap is satisfied $(t_{\varepsilon}(\theta) < 0)$, where only the necessary condition for overlap is satisfied $(t_{\varepsilon}(\theta) = 0)$, or where neither is satisfied $(t_{\varepsilon}(\theta) > 0)$. However, in light of Theorems 5.2 and 5.3, it is clear that t_{ε} could still contain (potentially large) intervals where it is zero, and generally, regions where $t_{\varepsilon}(\theta) < 0$ holds will often be found *in between* such regions where t_{ε} is the constant zero. Thus, there is still cause for concern that approximating t_{ε} to find regions where it is negative may be difficult. As such, in the next section we introduce an additional nonnegative function to replace the portions of t_{ε} where it is the constant zero.

Remark 5.4. Recall that we added smoothing in the definitions of a_{ε} and b_{ε} by squaring the $\operatorname{Arg}(\cdot)$ terms, as they otherwise may grow like the square root function when they increase from zero (or vice versa), behavior which can be difficult and expensive to resolve via interpolation. While ℓ_{ε} can also exhibit similar non-Lipschitz behavior when it transitions to being negative (and possibly elsewhere when it is already negative), we have intentionally not smoothed this term. The reason is that once an angle θ is found such that $\ell_{\varepsilon}(\theta) < 0$, there is no need to continue building an interpolant approximation. This angle can immediately be used to compute new level-set points to restart optimization and improve (lower) the current estimate to $\operatorname{sep}^{\lambda}_{\lambda}(A, B)$.

6 Locally supporting rays of pseudospectra and our certificate function d_{ε}

In this section, we propose a new function with which we can replace the constant-zero portions of t_{ε} . However, we begin with the following general definitions, which are variations of the concept

of a supporting hyperplane in \mathbb{R}^n [BV04, Chapter 2.5.2] specialized to \mathbb{C} , and a pair of related theoretical results.

Definition 6.1. Given a connected set $\mathcal{A} \subset \mathbb{C}$, a line $\mathcal{L} \subset \mathbb{C}$ supports \mathcal{A} at a point $z \in bd(\mathcal{A}) \cap \mathcal{L}$ if \mathcal{A} lies completely in one of the closed half-planes defined by \mathcal{L} .

Definition 6.2. Given a set $\mathcal{B} \subset \mathbb{C}$, a line $\mathcal{L} \subset \mathbb{C}$ locally supports \mathcal{B} at a point $z \in bd(\mathcal{B}) \cap \mathcal{L}$ if line \mathcal{L} supports $\mathcal{A} \cap \mathcal{N}$ at z for some neighborhood \mathcal{N} about z, where \mathcal{A} is a connected component of \mathcal{B} . A ray \mathcal{R} locally supports \mathcal{B} at $z \in bd(\mathcal{B}) \cap int \mathcal{R}$ if the line \mathcal{L} containing \mathcal{R} locally supports \mathcal{B} at z.

Note that if θ is a point where a_{ε} transitions from positive to zero (or vice versa), this implies that the ray \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A)$. Similarly, if θ is a point where b_{ε} transitions from positive to zero (or vice versa), then \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(B)$. Thus, it follows that if θ is a point where $a_{\varepsilon}+b_{\varepsilon}$ transitions from positive to zero (or vice versa), then \mathcal{R}_{θ} locally supports either $\Lambda_{\varepsilon}(A)$ or $\Lambda_{\varepsilon}(B)$ or both simultaneously (though not necessarily at the same point). Also note that if ℓ_{ε} transitions from zero to negative (or vice versa) at θ , then \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B)$. We now derive necessary conditions based on the eigenvalues of C_{θ} and S_{θ} for these scenarios. We first consider the case when \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A)$. Note that [BLO03, p. 371–373] also informally touches upon this subject and related issues for the specific case of vertical lines.

Lemma 6.3. Let $A \in \mathbb{C}^{m \times m}$, $\varepsilon \geq 0$, $z_0 \in \mathbb{C}$, $\theta \in \mathbb{R}$, and \mathcal{R}_{θ} be the ray defined in (4.1). If \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A)$, then the matrix C_{θ} defined in (4.3) has $\mathbf{i}\hat{r}$ with $\hat{r} > 0$ as a repeated eigenvalue with even algebraic multiplicity.

Proof. Without loss of generality, assume that $z_0 = 0$ and $\theta = 0$, and suppose that \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A)$ at $\hat{r} > 0$. Thus, $\hat{r} \in \operatorname{bd} \Lambda_{\varepsilon}(A)$, and so $\sigma_{\min}(A - \hat{r}I) = \varepsilon$ and $i\hat{r} \in \Lambda(C_{\theta})$ by Lemma 4.1. By Definition 6.2, there exists a neighborhood \mathcal{N} (in the open right half-plane) about \hat{r} such that $(\Lambda_{\varepsilon}(A) \cap \mathcal{N}) \setminus \mathcal{R}_{\theta}$ is connected. As \mathcal{R}_{θ} separates \mathcal{N} into $\mathcal{N}_1 = \{z \in \mathcal{N} : \text{Im } z > 0\}$ and $\mathcal{N}_2 = \{z \in \mathcal{N} : \text{Im } z < 0\}$, either $\Lambda_{\varepsilon}(A) \cap \mathcal{N}_1$ or $\Lambda_{\varepsilon}(A) \cap \mathcal{N}_2$ must be empty. Without loss of generality, suppose that $\Lambda_{\varepsilon}(A) \cap \mathcal{N}_1 = \emptyset$, and now consider how eigenvalue $i\hat{r}$ evolves as θ is varied, i.e., $\lambda(\theta) \in \Lambda(C_{\theta})$ with $\lambda(0) = \mathbf{i}\hat{r}$. By continuity, eigenvalue $\lambda(\theta)$ can either move up or down on the imaginary axis or it can move off the imaginary axis as the value of θ is increased from zero. If it moves along the imaginary axis, then locally, we have that $\lambda(\theta) = \mathbf{i}r(\theta)$, where $r: \mathbb{R} \to \mathbb{R}$ is continuous and $r(0) = \tilde{r}$. Since $\tilde{r} > 0$, there exists a $\theta_{\rm p} > 0$ such that $r(\theta) > 0$ for all $\theta \in (0, \theta_{\rm p})$. By Lemma 4.1, it thus follows that $r(\theta)e^{i\theta} \in \Lambda_{\varepsilon}(A)$ for all $\theta \in (0, \theta_p)$, but this contradicts the assumption that $\Lambda_{\varepsilon}(A) \cap \mathcal{N}_1$ is empty. Thus, $\lambda(\theta)$ must move off the imaginary axis as the value of θ is increased from zero. Since the eigenvalues of the Hamiltonian matrix C_{θ} are symmetric with respect to the imaginary axis, by continuity at least one pair of eigenvalues (or possibly more pairs non-generically) must coalesce on the imaginary axis at $i\hat{r}$ as $\theta \to 0$.

Now consider the case when \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B)$, which can happen at a boundary point of either $\Lambda_{\varepsilon}(A)$ or $\Lambda_{\varepsilon}(B)$, or a shared boundary point of both. Building on Lemma 6.3, we have the following result.

Lemma 6.4. Let $A \in \mathbb{C}^{m \times m}$, $B \in \mathbb{C}^{n \times n}$, $\varepsilon \geq 0$, $z_0 \in \mathbb{C}$, $\theta \in \mathbb{R}$, and \mathcal{R}_{θ} be the ray defined in (4.1). Furthermore, for matrix A, let C_{θ} be the matrix defined in (4.3), and let S_{θ} be its analogue for matrix B. If \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B)$ at a point $z \in \mathbb{C}$, then at least one, and possibly all, of the following conditions must hold:

- (i) C_{θ} and/or S_{θ} has $\mathbf{i}\hat{r}$ with $\hat{r} > 0$ as a repeated eigenvalue with even algebraic multiplicity,
- (ii) C_{θ} and S_{θ} have an eigenvalue $\mathbf{i}\hat{r}$ with $\hat{r} > 0$ in common.

Proof. Without loss of generality, we can assume that $z_0 = 0$ and $\theta = 0$, and so z is on the positive part of the real axis, i.e., $z = \hat{r}$ for some $\hat{r} > 0$. If \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B)$ at \hat{r} , either $\hat{r} \in \text{bd } \Lambda_{\varepsilon}(A)$ but not $\text{bd } \Lambda_{\varepsilon}(B)$ (or vice versa) or \hat{r} is a shared boundary point of both $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$. If \hat{r} is not a shared boundary point, then \mathcal{R}_{θ} must locally support either $\Lambda_{\varepsilon}(A)$ or $\Lambda_{\varepsilon}(B)$ at \hat{r} , and so Lemma 6.3 applies, yielding the "or" part of (i). Now suppose \hat{r} is a shared

boundary point, and so $\sigma_{\min}(A - \hat{r}I) = \sigma_{\min}(B - \hat{r}I) = \varepsilon$. Then by Lemma 4.1, $\mathbf{i}\hat{r}$ is an eigenvalue of both C_{θ} and S_{θ} , yielding (ii). Furthermore, \mathcal{R}_{θ} may or may not also locally support $\Lambda_{\varepsilon}(A)$ and/or $\Lambda_{\varepsilon}(B)$ at \hat{r} . All four scenarios are possible, with the "and" part of (i) corresponding to when the ray simultaneously locally supports both $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ at \hat{r} .

Recall the set of roots $\mathcal{T}_{\varepsilon}$ of t_{ε} , which is defined in (5.2b). If $\theta \in \mathcal{T}_{\varepsilon}$, then the necessary condition for overlap $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = 0$ is satisfied, and so \mathcal{R}_{θ} intersects both $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$. However, as $\ell_{\varepsilon}(\theta) = 0$, the sufficient condition is not met, and via Theorem 5.3, it follows that $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ either have no points in common along \mathcal{R}_{θ} , or at most only boundary points in common. For a function to replace the regions of t_{ε} where $t_{\varepsilon}(\theta) = 0$, i.e., $\mathcal{T}_{\varepsilon}$, we propose a function $d_{\varepsilon}^{AB} : \mathcal{T}_{\varepsilon} \to [0, \infty)$ that is a measure of how close $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ are to sharing a boundary point along \mathcal{R}_{θ} . To that end, let

$$d_{\varepsilon}^{AB}(\theta) \coloneqq \min\{d_{\varepsilon}^{A}(\theta), d_{\varepsilon}^{B}(\theta)\}, \quad \text{where}$$
(6.1a)

$$d_{\varepsilon}^{A}(\theta) \coloneqq \min\{f_{A}(r,\theta) - \varepsilon : z_{0} + re^{\mathbf{i}\theta} \in \mathcal{R}_{\theta} \cap \operatorname{bd}\Lambda_{\varepsilon}(B)\},$$
(6.1b)

$$d^B_{\varepsilon}(\theta) \coloneqq \min\{f_B(r,\theta) - \varepsilon : z_0 + r e^{\mathbf{i}\theta} \in \mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon}(A)\},$$
(6.1c)

where f_A is defined in (4.2) for matrix A and f_B is its analogue for matrix B. Since $\theta \in \mathcal{T}_{\varepsilon}$, both $\mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon}(A)$ and $\mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon}(B)$ must be nonempty, and so the functions are well defined. The purpose of d_{ε}^A is to provide a nonnegative measure of how close $\Lambda_{\varepsilon}(B)$ is to touching $\Lambda_{\varepsilon}(A)$ along the given ray \mathcal{R}_{θ} , and vice versa for d_{ε}^B . Note that if $\mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon}(A) \cap \operatorname{bd} \Lambda_{\varepsilon}(B) \neq \emptyset$, then $d_{\varepsilon}^A(\theta) = d_{\varepsilon}^B(\theta) = 0$, but otherwise $d_{\varepsilon}^A(\theta)$ and $d_{\varepsilon}^B(\theta)$ are typically not the same value. While technically d_{ε}^A alone (or d_{ε}^B) would suffice as a closeness measure of the two pseudospectra along a given ray, we have observed that their pointwise minimum, i.e., d_{ε}^{AB} , is often cheaper to approximate. Important properties of d_{ε}^{AB} are summarized in the following statement.

Theorem 6.5 (Properties of d_{ε}^{AB}). Let $A \in \mathbb{C}^{m \times m}$, $B \in \mathbb{C}^{n \times n}$, $\varepsilon \ge 0$, and $z_0 \in \mathbb{C}$ be such that $\varepsilon \ge 0$ is not a singular value of either $A - z_0I$ or $B - z_0I$, and let \mathcal{R}_{θ} be the ray defined in (4.1). Furthermore, let d_{ε}^{AB} be as defined in (6.1) on domain $\mathcal{T}_{\varepsilon}$ defined in (5.2b). Then for any point $\theta \in \mathcal{T}_{\varepsilon}$, the following statements hold:

(i) $d_{\varepsilon}^{AB}(\theta) \ge 0$,

(ii)
$$d_{\varepsilon}^{AB}(\theta) = 0 \quad \iff \quad \mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B) \neq \varnothing \quad \iff \quad \mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon}(A) \cap \operatorname{bd} \Lambda_{\varepsilon}(B) \neq \varnothing,$$

- (iii) d_{ε}^{AB} is continuous at θ if every eigenvalue ir, of either C_{θ} or S_{θ} , that attains the minimum in $d_{\varepsilon}^{AB}(\theta)$ is simple,
- (iv) d_{ε}^{AB} is differentiable at θ if there are no ties for $d_{\varepsilon}^{AB}(\theta)$, i.e., it is attained via $f_A(r,\theta)$ or $f_B(r,\theta)$ but not both, the corresponding minimum singular value is simple, and there is a single eigenvalue ir, of either C_{θ} or S_{θ} as appropriate, that attains $d_{\varepsilon}^{AB}(\theta)$, where this eigenvalue is simple.

Proof. Statements (i) and (ii) are simple but important direct consequences of the definition of d_{ε}^{AB} and the fact that its domain is restricted to $\mathcal{T}_{\varepsilon}$, since otherwise $d_{\varepsilon}^{AB}(\theta)$ could be negative (or undefined) for some θ and the equivalences in (ii) would not hold. For statement (iii), consider $d_{\varepsilon}^{B}(\theta)$ and recall that by Lemma 4.1, $z_{0} + \hat{r}e^{i\theta} \in \mathcal{R}_{\theta} \cap \mathrm{bd} \Lambda_{\varepsilon}(A)$ is always associated with an eigenvalue $\mathbf{i}\hat{r}$ of C_{θ} . Since eigenvalues are continuous, eigenvalue $\mathbf{i}\hat{r}$ can either move continuously along the positive portion of the imaginary axis or leave this region as θ is varied. Clearly, the former case cannot cause a discontinuity in d_{ε}^{B} , so consider the latter. By the assumption on ε , zero can never be an eigenvalue of C_{θ} for any θ , and clearly the eigenvalues of a matrix are all finite. Thus, if an eigenvalue leaves the positive portion of the imaginary axis, it cannot be by going through the origin or infinity. Since the eigenvalues of the Hamilton matrix C_{θ} are symmetric with respect to the imaginary axis, a simple eigenvalue cannot leave the imaginary axis, and a repeated eigenvalue is excluded by assumption; hence, d_{ε}^{B} must be continuous at θ . The same argument shows that d_{ε}^{A} is continuous at θ under the analogous assumptions for the eigenvalues of S_{θ} , and so d_{ε}^{AB} is continuous at θ . For (iv), the assumptions mean that there are no ties for the min functions and standard perturbation theory for simple singular values and simple eigenvalues applies.



Figure 6.1: For two randomly generated matrices $A, B \in \mathbb{C}^{10 \times 10}$, the left pane shows their eigenvalues (respectively x's and dots), and $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$ (respectively solid and dotted contours) for $\varepsilon = 0.3 > \operatorname{sep}_{\lambda}^{D}(A, B)$. The search point z_{0} is the origin; rays emanating from it are depicted by dashed lines. The right pane shows a corresponding plot of d_{ε} , where its components are plotted as follows: $a_{\varepsilon} + b_{\varepsilon}$ (dotted), ℓ_{ε} (dashed), and d_{ε}^{AB} (solid). For $\theta = -\frac{1}{2}\pi$, it can be seen in the left pane that \mathcal{R}_{θ} only passes through $\Lambda_{\varepsilon}(B)$ and so $d_{\varepsilon}(\theta) = a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) > 0$ in the right pane. Meanwhile for $\theta = \frac{1}{2}\pi$, \mathcal{R}_{θ} passes through $\inf \Lambda_{\varepsilon}(A) \cap \inf \Lambda_{\varepsilon}(B)$ and so $d_{\varepsilon}(\theta) = \ell_{\varepsilon}(\theta) < 0$. Finally, for $\theta = 0$, while \mathcal{R}_{θ} passes through both $\Lambda_{\varepsilon}(A)$ and $\Lambda_{\varepsilon}(B)$, it never does so simultaneously, hence $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = \ell_{\varepsilon}(\theta) = 0$ and $d_{\varepsilon}(\theta) = d_{\varepsilon}^{AB}(\theta) > 0$.

While Theorem 6.5 verifies that d_{ε}^{AB} is reasonably well behaved, d_{ε}^{AB} may have jump discontinuities. However, d_{ε}^{AB} is discontinuous at point $\theta \in \operatorname{int} \mathcal{T}_{\varepsilon}$ only if two conditions simultaneously hold: \mathcal{R}_{θ} locally supports $\Lambda_{\varepsilon}(A)$ or $\Lambda_{\varepsilon}(B)$ at a point $z_0 + \hat{r}e^{i\theta}$ with $\hat{r} > 0$, and this value \hat{r} is the one that attains the value of $d_{\varepsilon}^{AB}(\theta)$. As a result, we expect such discontinuities to be relatively few, and so this should not be a problem in practice. Functions d_{ε}^A and d_{ε}^B typically do not have non-Lipschitz behavior when they transitions to/from zero, and so we have not added smoothing when them and defining d_{ε}^{AB} . When f^{D} has a unique minimizer, d_{ε}^{AB} only has a single root for $\varepsilon = \operatorname{sep}_{\lambda}^{D}(A, B)$.

Combining our three constituent pieces, we now define $d_{\varepsilon} : (-\pi, \pi] \to \mathbb{R}$, our key function for our interpolation-based globality certificate for $\operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$:

$$d_{\varepsilon}(\theta) \coloneqq \begin{cases} a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) & \text{if } a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) > 0, \\ \ell_{\varepsilon}(\theta) & \text{if } \ell_{\varepsilon}(\theta) < 0, \\ d_{\varepsilon}^{AB}(\theta) & \text{otherwise.} \end{cases}$$
(6.2)

In Fig. 6.1, we plot d_{ε} for a sample problem with $\varepsilon > \operatorname{sep}_{\lambda}^{D}(A, B)$ in order to illustrate the different components of d_{ε} . Recalling that $a_{\varepsilon} + b_{\varepsilon}$ is a nonnegative function and so is d_{ε}^{AB} on its domain, we immediately have the following global convergence conditions as a corollary of Theorems 5.2, 5.3, and 6.5.

Corollary 6.6 (Global convergence for $\operatorname{sep}_{\lambda}^{D}(A, B)$ via d_{ε}). Let $A \in \mathbb{C}^{m \times m}$, $B \in \mathbb{C}^{n \times n}$, $\varepsilon \geq 0$, and $z_0 \in \mathbb{C}$ be such that ε is not a singular value of either $A - z_0I$ or $B - z_0I$, and let d_{ε} be the function defined in (6.2). Then

$$\min_{\theta \in (-\pi,\pi]} d_{\varepsilon}(\theta) < 0 \quad \Longleftrightarrow \quad \mu(\{\theta \in (-\pi,\pi] : d_{\varepsilon}(\theta) < 0\}) > 0 \quad \Longleftrightarrow \quad \varepsilon > \operatorname{sep}_{\lambda}^{\mathrm{D}}(A,B)$$

In the process of devising d_{ε} , we considered many different possibilities but found that these alternatives were significantly more expensive to use than d_{ε} , even if they had fewer jumps or

even none. For example, we considered an entirely continuous alternative to d_{ε} that replaced its d_{s}^{AB} portions with a continuous measure of the distance to any of the necessary conditions in Lemma 6.4 holding. However, this function often had more complicated behavior and many many roots than d_{ε} because the necessary conditions in Lemma 6.4 hold for any θ such that \mathcal{R}_{θ} locally supports either of the two pseudospectra or their intersection, and possibly at other angles as well. Even when incorporating smoothing to address non-Lipschitz behavior at roots, this alternative was still much more expensive to approximate than d_{ε} . We also tried replacing d_{ε}^{AB} with $\min\{\mu(\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A)), \mu(\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B))\}\$ and other continuous alternatives, although these choices still resulted in jumps when combined when used in conjunction with $a_{\varepsilon} + b_{\varepsilon}$ and ℓ_{ε} . But these choices were more expensive to approximate than d_{ε}^{AB} because they generally had more complicated behaviors than d_{ε}^{AB} , e.g., more nonsmooth points, more oscillatory behavior, etc. Finally, we considered just using the smallest pairwise distance between points in $\mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon}(A)$ and $\mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon}(B)$. This is quite similar to d_{ε}^{AB} and can have similar discontinuities, but it too ended up being more expensive to approximate than d_{ε}^{AB} . That all said, none of the alternatives we considered were prohibitively expensive; using any of them to compute $\operatorname{sep}_{\lambda}^{D}(A, B)$ was still much faster than the method of Gu and Overton, even though they were generally not as fast as our ultimate choice for d_{ε} .

Remark 6.7. Another approach to computing $sep_{\lambda}^{D}(A, B)$ is via

$$s(\theta) \coloneqq \min_{r \in \mathbb{R}} s_{\theta}(r), \quad where \quad s_{\theta}(r) \coloneqq \max\{\sigma_{\min}(F_A(r,\theta)), \sigma_{\min}(F_B(r,\theta))\}, \tag{6.3}$$

i.e., $s(\theta)$ is the minimal value f^{D} takes along the line defined by θ and passing through some $z_{0} \in \mathbb{C}$. It is then immediate that

$$\operatorname{sep}_{\lambda}^{\mathrm{D}}(A,B) = \min_{\theta \in [0,\pi)} s(\theta), \tag{6.4}$$

as this simply rewrites (1.3b) in polar coordinates about z_0 . Thus, using Chebfun to approximate s and then find a global minimizer in $[0,\pi)$ provides another way to obtain sep¹₁(A,B). One drawback of this approach is that for any given θ , evaluating $s(\theta)$ is much more expensive than evaluating $d_{\varepsilon}(\theta)$. As we explain in detail in the next section, evaluating $d_{\varepsilon}(\theta)$ is essentially direct, since it only requires solving two eigenvalue problems of order 2m and 2n and this is generally the dominant cost. Meanwhile, computing $s(\theta)$ involves finding a global minimizer of s_{θ} , which requires iteration. Although we can use Lemma 4.1 to construct such an iteration, similar to the level-set methods of [BB90, BS90] for computing the \mathcal{H}_{∞} norm, the resulting algorithm to compute $s(\theta)$ would generally only be linearly convergent; the key difference between here and the \mathcal{H}_{∞} -norm setting is that s_{θ} , due to being a max of two min functions, will generally will be nonsmooth at its minimizers. Consequently, evaluating $s(\theta)$ would require solving multiple eigenvalue problems of 2m and 2n. Another issue is that although s is continuous, it is still nonsmooth, and it is generally more expensive for Chebfun to detect nonsmooth points than jumps; see [PPT09, Tre20]. Finally, a third downside is that using Chebfun to precisely compute a (likely unique) global minimizer of some function, e.g., s, is a significantly more numerically challenging task than what we ask of Chebfun inside our algorithm using d_{ε} , i.e., to find any point where d_{ε} is negative, since as we have shown, the set of such points has positive measure when $\varepsilon > \sup_{\lambda}^{D}(A, B)$. Thus, when attempting to compute $\operatorname{sep}_{\lambda}^{D}(A, B)$ by applying Chebfun to (6.4), we nevertheless recommend subsequently refining its computed result by applying local optimization to f^{D} initialized from the point in the complex plane found by Chebfun.

7 Implementation and the cost of our method

We now discuss how to implement our $\operatorname{sep}_{\lambda}^{D}(A, B)$ algorithm, which we have done in MATLAB, and describe its overall work complexity. We give detailed remarks in the following subsections, while high-level pseudocode is given in Algorithm 7.1.

7.1 Choosing a search point

Regarding what search point z_0 to use, we recommend the average of all the distinct eigenvalues of A and B. This helps to ensure the whole domain of $d_{\varepsilon}(\theta)$ is relevant. Otherwise, if for a given

Algorithm 7.1 Interpolation-based Globality Certificate Algorithm for $\operatorname{sep}_{\lambda}^{D}(A, B)$

Input: $A \in \mathbb{C}^{m \times m}$, $B \in \mathbb{C}^{n \times n}$, "search point" $z_0 \in \mathbb{C}$, and $z_{\text{init}} \in \mathbb{C}$. **Output:** $\varepsilon \approx \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$.

```
1: while true do
```

 $\varepsilon \leftarrow$ computed locally/globally minimal value of $f^{\rm D}$ initialized from $z_{\rm init}$ 2: // Begin approximating d_{ε} to assert convergence or find new starting points 3: 4: $p_{\varepsilon} \leftarrow 1$ // Initial guess for polynomial interpolant p_{ε} for approximating d_{ε} while p_{ε} does not sufficiently approximate d_{ε} do 5: $[\theta_1, \ldots, \theta_q] \leftarrow$ new sample points from $(-\pi, \pi]$ 6: // If new starting points are detected, restart optimization to lower ε : 7: if $d_{\varepsilon}(\theta_j) < 0$ for some $j \in \{1, \ldots, q\}$ then 8: $z_{\text{init}} \leftarrow \text{a point in bd} \{ \mathcal{R}_{\theta_i} \cap \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B) \} \setminus \{ z_0 \}$ 9: goto line 2 // Restart optimization from z_{init} 10: end if 11: // Otherwise, no starting points detected, keep improving p_{ε} : 12: $p_{\varepsilon} \leftarrow \text{improved polynomial interpolant of } d_{\varepsilon} \text{ via } \theta_1, \ldots, \theta_q$ 13:14: end while // p_{ε} approximates d_{ε} well and no new starting points were encountered 15:// However, do a final check before asserting that d_{ε} is nonnegative: 16: $[\theta_1, \ldots, \theta_q] = \arg\min p_{\varepsilon}(\theta)$ 17:if $d_{\varepsilon}(\theta_j) < 0$ for some $j \in \{1, \ldots, q\}$ then 18: $z_{\text{init}} \leftarrow \text{a point in bd} \{ \mathcal{R}_{\theta_i} \cap \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B) \} \setminus \{ z_0 \}$ 19:goto line 2 // Restart optimization from z_{init} 20:21:else return // $p_{\varepsilon} \approx d_{\varepsilon}$ and $\implies \varepsilon \approx \operatorname{sep}_{\lambda}^{\mathrm{D}}(A, B)$ 22: end if 23:24: end while

value of ε , z_0 is chosen far from the pseudospectra of A and B, then $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = 0$ would only hold on a very small subset of $(-\pi, \pi]$, which in turn would likely make it harder to find the regions where $d_{\varepsilon}(\theta)$ is negative. On every round, our code checks that the choice of z_0 still satisfies our needed assumptions and perturbs it slightly if it does not (in practice, we have not observed that this is necessary). Finally, if the pseudospectra of A and B both have real-axis symmetry, by choosing z_0 on the real axis, it is then only necessary to approximate $d_{\varepsilon}(\theta)$ on $[0, \pi]$.

7.2 Evaluating $d_{\varepsilon}(\theta)$ and its cost

Given some θ , evaluating $d_{\varepsilon}(\theta)$ proceeds as follows. First, the eigenvalues of both C_{θ} and S_{θ} are computed. For increased reliability, it is recommended that this be done via a structurepreserving eigensolver such as [BMX99]. From these spectra, it is then trivial to calculate the value of $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta)$ via (4.5a). If $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) > 0$, then the value of $d_{\varepsilon}(\theta)$ has been computed. Otherwise, evaluating $d_{\varepsilon}(\theta)$ requires the following additional computations, which begins with obtaining the value of $\ell_{\varepsilon}(\theta)$. To that end, we compute $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A)$ and $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B)$. Considering the former, we want to determine the values r > 0 such that $f_A(r, \theta) = \varepsilon$, and via Lemma 4.1, we have the following sorted list of candidate values $0 = r_0 < r_1 < \ldots < r_q$ that may satisfy this equality, where $\mathbf{i}r_j$ for $j = 1, \ldots, q$ are eigenvalues of C_{θ} and we have added $r_0 = 0$. Then to compute $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A)$, we must assert which intervals on \mathcal{R}_{θ} , defined by $[r_{j-1}, r_j]$ for $j = 1, \ldots, q$,

NOTE: To keep the pseudocode a reasonable length, we make some simplifying assumptions: optimization converges to local/global minimizers exactly, z_{init} computed in lines 9 and 19, for restarting optimization, is never a stationary point of f^{D} , and the "search point" z_0 is such that all encountered values of ε are not singular values of $\sigma_{\min}(A - z_0I)$ and $\sigma_{\min}(B - z_0I)$, per the assumptions given in §4 and §5. Lines 3-15 describe the core of the interpolation-based globality certificate, where we only give a broad outline of the interpolation process for approximating d_{ε} ; note that for numerical reasons, each certificate should actually be done with $\tilde{\varepsilon} = (1 - \tau)\varepsilon$, where $\tau \in (0, 1)$ is some relative tolerance. See §7.2 and §7.3 for more implementation details.

are also in $\Lambda_{\varepsilon}(A)$. There are several ways to do this but a simple and robust way is to just evaluate $f_A(\hat{r}_j, \theta)$ for $\hat{r}_j = 0.5(r_{j-1} + r_j)$ over $j = 1, \ldots, q$; since $f_A(\hat{r}_j, \theta) \neq \varepsilon$, the corresponding interval is not in $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A)$ if and only if $f_A(\hat{r}_j, \theta) > \varepsilon$. Note that it does not matter if we have two or more adjacent intervals in our computed version of $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A)$. An analogous computation yields $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(B)$. With these two sets computed, calculating the amount of their overlap along the given ray, i.e., $-\ell_{\varepsilon}(\theta)$, is straightforward. If $\ell_{\varepsilon}(\theta) < 0$, then the evaluation of $d_{\varepsilon}(\theta)$ is done and the boundary points of $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B)$ have been also been computed, which are used to restart optimization. However, if $\ell_{\varepsilon}(\theta) = 0$, then finally we must compute $d_{\varepsilon}^{AB}(\theta)$ in order to complete the computation of $d(\theta)$, though this is this is straightforward to do from the definition of $d_{\varepsilon}^{AB}(\theta)$ given in (6.1) and the previous computations.

Recalling our assumption that $m \leq n$, evaluating $d_{\varepsilon}(\theta)$ is $\mathcal{O}(n^3)$ work if done in the following manner. Computing all of the eigenvalues of C_{θ} and S_{θ} is $\mathcal{O}(n^3)$ work, and that is all there is to do when $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) > 0$. But when $a_{\varepsilon}(\theta) + b_{\varepsilon}(\theta) = 0$, computing $d_{\varepsilon}(\theta)$ additionally requires computing the values of $f_A(r,\theta)$ and $f_B(r,\theta)$ for different values of r. While the number of values of r is often only a handful, in the worst case, it can be $\mathcal{O}(m+n)$. Hence, if we were to evaluate this pair of functions by computing SVDs, we would exceed the stated $\mathcal{O}(n^3)$ work complexity bound by a factor of n. Fortunately, there is a more efficient option due to Lui for fast plotting of pseudospectra [Lui97]. Since A is square, it has a Schur decomposition $A = UTU^*$, where U is unitary and T is triangular, and moreover, since unitary transformations do not alter the pseudospectrum, $\Lambda_{\varepsilon}(A) = \Lambda_{\varepsilon}(T)$ holds. The key benefit of this transformation is that at any point $z_0 + re^{i\theta} \in \mathbb{C}$, we have that $T - (z_0 + re^{i\theta})I$ remains in triangular form, and so inverse iteration can be done to compute this shifted matrix's minimum singular value using backsolves that only require quadratic work as opposed to the usual cubic work for solving a linear system. We need only compute and store Schur decompositions of A and B once in an offline phase, which is cubic work, and then we can evaluate $f_A(r,\theta)$ and $f_B(r,\theta)$ for any r and θ in a most $\mathcal{O}(n^2)$ work under the mild assumption that inverse iteration converges in relatively few steps.⁴ Hence, evaluating $d_{\varepsilon}(\theta)$ can always be done within $\mathcal{O}(n^3)$ work. In our own experience, we have seen that ten iterations is generally more than sufficient to compute $f_A(r,\theta)$ and $f_B(r,\theta)$ accurately to the full precision of the hardware, and that this technique is already faster than computing the full SVD for matrices as small as 50×50 .

7.3 Approximating d_{ε} and restarting

To approximate d_{ε} , we use Chebfun, as it is rather adept at approximating functions with nonsmooth points and/or discontinuities. As Chebfun normally provides groups of points to evaluate simultaneously (line 6 of Algorithm 7.1), these evaluations of d_{ε} can be done in parallel; see [Mit21, Section 5.2] for more details. Furthermore, if $d_{\varepsilon}(\theta) < 0$ for any of current group of points provided by Chebfun, we immediately halt Chebfun and use the detected boundary points of $\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon}(A) \cap \Lambda_{\varepsilon}(B)$ (except for z_0) to restart optimization (lines 7–11 of Algorithm 7.1). This is accomplished by throwing an error when a point is encountered such that $d_{\varepsilon}(\theta) < 0$ holds, which causes Chebfun to be aborted. By subsequently catching our own thrown error, we can resume our program to restart another round of optimization.

7.4 Finding minimizers

Like many other optimization-with-restarts algorithms, it will be necessary to use a monotonic optimization solver, i.e., one that always decreases the objective function on every iteration, which is the case for most unconstrained optimization solvers. Minimizers of $f^{\rm D}$ will almost always be nonsmooth, and at best, we can expect linear convergence from a nonsmooth optimization solver. However, since there are only two real variables, we expect the number of iterations needed to converge to be relatively small. Thus, as evaluating $f^{\rm D}$ and its gradient is significantly cheaper than evaluating d_{ε} , and we expect far fewer function evaluations for the former than the latter, the cost of Algorithm 7.1 will generally not be dominated by the optimization phases.

⁴For more details on the actual inverse-iteration-based algorithm, including pseudocode and code examples, see [Lui97] and [TE05, Chapter 39], but note that the latter has the following typo: In "Core EigTool algorithm" [TE05, p. 375], the second to last line should be sigmin(j,k) = 1/sqrt(sig);, not sigmin(j,k) = sqrt(sig);.

To find minimizers of f^{D} using only gradient information, we use GRANSO: GRadient-based Algorithm for Non-Smooth Optimization [Mita]. GRANSO implements the BFGS-SQP nonsmooth optimization algorithm of [CMO17], which can handle nonsmooth constraints, but for problems without constraints, it reduces to BFGS with the line search of [LO13], a combination which Lewis and Overton have studied and advocated as a method for nonsmooth optimization. While there are no convergence results for BFGS for general nonsmooth optimization, it nevertheless seems to reliably and accurately converge to nonsmooth stationary values. Indeed, in their concluding remarks [LO13, p. 160], Lewis and Overton wrote "In our experience with functions with bounded sublevel sets, BFGS essentially always generates function values converging linearly to a Clarke stationary value, with exceptions only in cases that we attribute to the limits of machine precision. We speculate that, for some broad class of reasonably well-behaved functions, this behavior is almost sure." Since f^{D} is locally Lipschitz as long as $\operatorname{sep}_{\lambda}^{D}(A, B) > 0$ and has bounded level sets, we expect that BFGS will also be an efficient and reliable tool in our setting. For improved theoretical guarantees, one could follow up optimization via BFGS with a phase of the gradient sampling algorithm [BLO05], which would ensure convergence to nonsmooth stationary values of $f^{\rm D}$ when $\operatorname{sep}_{\lambda}^{\rm D}(A, B) > 0$. However, for simplicity, we only use BFGS here.

When restarting optimization, our certificate may provide many new starting points. Restarting from just one would give the smallest chance of converging to a global minimizer on this round, while restarting from them all could be a waste of time, particularly if this ends up just returning the same minimizer over and over again. In practice, one could prioritize them in terms of most promising first and limit the total number used. On multi-core machines, optimization can be run from multiple starting points in parallel.

7.5 Terminating the algorithm

In addition to the convergence tests described in Algorithm 7.1, it is also necessary to terminate the algorithm if consecutive estimates for $\operatorname{sep}_{\lambda}^{D}(A, B)$ are identical. The reason is that we cannot expect optimization solvers to find minimizers *exactly*. If a global minimizer \tilde{z} is obtained only up to some rounding error, then $\operatorname{sep}_{\lambda}^{D}(A, B)$ has essentially been computed, but our certificate may still detect that the algorithm has not *truly* converged to a global minimizer, and in this case, the algorithm may try to restart optimization (unsuccessfully). This is also part of the reason why the certificates should actually be performed with $\tilde{\varepsilon} = (1 - \tau)\varepsilon$, as described in the note under Algorithm 7.1.

7.6 The overall work complexity and using lines instead of rays

In the worst case, the overall work complexity to perform the interpolation-based globality certificates is $\mathcal{O}(kn^3)$, where k is the total number of function evaluations (over all values of ε encountered). As restarts tend to happen quickly, k is roughly equal to the number of evaluations needed to approximate d_{ε} when $\varepsilon = \operatorname{sep}_{\lambda}^{D}(A, B)$, and as we will see in the numerical experiments, k can generally be considered to be like a large constant, although it is influenced by the geometry of the two pseudospectra.

When implementing the algorithm, the definition of d_{ε} can be modified so that it considers lines through z_0 instead of rays emanating from z_0 . This can be beneficial, since we always get information for the direction $\theta + \pi$ when considering \mathcal{R}_{θ} , and so this modified version of d_{ε} need only be interpolated on $[0, \pi]$. Function a_{ε} measures the minimum argument of $-\mathbf{i}\lambda$ over each eigenvalue λ of C_{θ} , so when using lines instead of rays, it must also consider the minimum angle with respect to the negative real axis. These additional angles are computed by simply switching the sign of the imaginary part of each eigenvalue λ . The same change is made for b_{ε} , while modifying ℓ_{ε} and d_{ε}^{AB} is straightforward. While using lines often results in less overall work, this is not always the case, as it can sometimes make d_{ε} more complicated and thus more expensive to approximate.

8 Algorithms for $sep_{\lambda}^{V}(A, B)$

We now briefly turn to the problem of computing Varah's $\operatorname{sep}_{\lambda}(A, B)$. We first answer whether or not Algorithm 7.1 extends to $\operatorname{sep}_{\lambda}^{V}(A, B)$ and then propose a different algorithm to compute $\operatorname{sep}_{\lambda}^{V}(A, B)$.

8.1 Does Algorithm 7.1 extend to $sep_{\lambda}^{V}(A, B)$?

In the construction of function d_{ε} for computing $\operatorname{sep}_{\lambda}^{D}(A, B)$, nowhere have we needed that the same value of ε be used for the pseudospectra of A and B. Thus for Varah's version of $\operatorname{sep}_{\lambda}(A, B)$, we can analogously define

$$d_{\varepsilon_1,\varepsilon_2}(\theta) \coloneqq \begin{cases} a_{\varepsilon_1}(\theta) + b_{\varepsilon_2}(\theta) & \text{if } a_{\varepsilon_1}(\theta) + b_{\varepsilon_2}(\theta) > 0, \\ \ell_{\varepsilon_1,\varepsilon_2}(\theta) & \text{if } \ell_{\varepsilon_1,\varepsilon_2}(\theta) < 0, \\ d_{\varepsilon_1,\varepsilon_2}^{AB}(\theta) & \text{otherwise,} \end{cases}$$
(8.1)

where

$$\ell_{\varepsilon_{1},\varepsilon_{2}}(\theta) \coloneqq -\mu \left(\mathcal{R}_{\theta} \cap \Lambda_{\varepsilon_{1}}(A) \cap \Lambda_{\varepsilon_{2}}(B)\right), \\ d_{\varepsilon_{1},\varepsilon_{2}}^{AB}(\theta) \coloneqq \min\{d_{\varepsilon_{1},\varepsilon_{2}}^{A}(\theta), d_{\varepsilon_{1},\varepsilon_{2}}^{B}(\theta)\}, \\ d_{\varepsilon_{1},\varepsilon_{2}}^{A}(\theta) \coloneqq \min\{f_{A}(r,\theta) - \varepsilon_{1} : \mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon_{2}}(B)\}, \\ d_{\varepsilon_{1},\varepsilon_{2}}^{B}(\theta) \coloneqq \min\{f_{B}(r,\theta) - \varepsilon_{2} : \mathcal{R}_{\theta} \cap \operatorname{bd} \Lambda_{\varepsilon_{1}}(A)\},$$

and f_A is defined in (4.2) for matrix A, while f_B is its analogue for matrix B. Although this will not allow us to compute $\sup_{\lambda}^{V}(A, B)$ to arbitrary accuracy, we do have the following necessary condition as another corollary of Theorem 5.3.

Corollary 8.1 (A necessary condition for $\varepsilon_1 + \varepsilon_2 = \sup_{\lambda}^{V}(A, B)$ via $d_{\varepsilon_1, \varepsilon_2}$). Let $A \in \mathbb{C}^{m \times m}$, $B \in \mathbb{C}^{n \times n}$, $\varepsilon_1, \varepsilon_2 \ge 0$, and $z_0 \in \mathbb{C}$ be such that ε_1 and ε_2 are, respectively, not singular values of $A - z_0I$ and $B - z_0I$, and let $d_{\varepsilon_1, \varepsilon_2}$ be the function defined in (8.1). Then

$$\min_{\theta \in (-\pi,\pi]} d_{\varepsilon_1,\varepsilon_2}(\theta) < 0 \quad \Longleftrightarrow \quad \mu(\{\theta \in (-\pi,\pi] : d_{\varepsilon_1,\varepsilon_2}(\theta) < 0\}) > 0,$$

and

$$\varepsilon_1 + \varepsilon_2 > \sup_{\lambda}^{\mathcal{V}}(A, B)$$
 if $\min_{\theta \in (-\pi, \pi]} d_{\varepsilon_1, \varepsilon_2}(\theta) < 0.$

As the last statement in Corollary 8.1 is not if-and-only-if, $d_{\varepsilon_1,\varepsilon_2}$ does not allow us compute $\operatorname{sep}_{\lambda}^{\mathrm{V}}(A, B)$ with guaranteed accuracy. However, by modifying Algorithm 7.1 to instead find minimizers of f^{V} and use $d_{\varepsilon_1,\varepsilon_2}$, we can compute locally optimal upper bounds for $\operatorname{sep}_{\lambda}^{\mathrm{V}}(A, B)$ that at least guarantee the necessary condition $\operatorname{int} \Lambda_{\varepsilon_1}(A) \cap \operatorname{int} \Lambda_{\varepsilon_2}(B) = \emptyset$ is satisfied, as this is equivalent to $\min_{\theta \in (-\pi,\pi]} d_{\varepsilon_1,\varepsilon_2}(\theta) = 0$. This is notably better than just computing upper bounds via finding minimizers of f^{V} , since the corresponding values of ε_1 and ε_2 associated with minimizers are not guaranteed to satisfy this necessary condition. However, when either $\varepsilon_1 = 0$ or $\varepsilon_2 = 0$ holds at the computed minimizer, note that $\operatorname{int} \Lambda_{\varepsilon_1}(A) = \emptyset$ or $\operatorname{int} \Lambda_{\varepsilon_2}(B) = \emptyset$ holds, and so satisfying the necessary condition does not preclude the possibility that an eigenvalue of A may be in $\operatorname{int} \Lambda_{\varepsilon_2}(B)$ or vice versa. Thus, when approximating $\operatorname{sep}_{\lambda}^{\mathrm{V}}(A, B)$ via this extended algorithm, one should always compute

$$\tilde{\varepsilon} = \min\left\{\min_{\lambda \in \Lambda(B)} \sigma_{\min}(A - \lambda I), \min_{\lambda \in \Lambda(A)} \sigma_{\min}(B - \lambda I)\right\},\tag{8.2}$$

which computes an upper bound $\tilde{\varepsilon} \geq \operatorname{sep}_{\lambda}^{V}(A, B)$ such that no eigenvalues of A are in the interior of $\operatorname{int} \Lambda_{\tilde{\varepsilon}}(B)$ and vice versa. Nevertheless, when optimization finds minimizers where neither ε_1 nor ε_2 is zero, then our certificate can be used to restart optimization if the necessary condition does not hold, and hence obtain a better estimate for $\operatorname{sep}_{\lambda}^{V}(A, B)$.

8.2 A different Chebfun-based algorithm to compute $sep_{\lambda}^{V}(A, B)$

Given $z_0 \in \mathbb{C}$, let function $v : [0, \pi) \to \mathbb{R}$ be defined as

$$v(\theta) \coloneqq \min_{r \in \mathbb{R}} v_{\theta}(r), \quad \text{where} \quad v_{\theta}(r) \coloneqq \sigma_{\min}(F_A(r,\theta)) + \sigma_{\min}(F_B(r,\theta)), \tag{8.3}$$

i.e., $v(\theta)$ is the minimal value f^{V} takes along the line defined by θ and passing through z_0 . It then immediately follows that

$$\operatorname{sep}_{\lambda}^{\mathcal{V}}(A,B) = \min_{\theta \in [0,\pi)} v(\theta).$$
(8.4)

Since v is continuous function defined on a finite interval, as in the alternative $\operatorname{sep}_{\lambda}^{D}(A, B)$ algorithm discussed in Remark 6.7, we can consider approximating v with Chebfun in order to solve (8.4).

Unfortunately evaluating v for a given θ is quite difficult, as the level-set iteration for finding a global minimizer of s_{θ} described in Remark 6.7 does not extend to v_{θ} . However, for some $\varepsilon > \sup_{\lambda}^{V}(A, B)$, say, $\varepsilon = f^{V}(z_{0})$, we can easily calculate a finite interval $[r_{1}, r_{2}]$ such that $v_{\theta}(r) > \varepsilon$ must hold for all $r \notin [r_{1}, r_{2}]$. To do this, we simple apply Lemma 4.1 to obtain the two extremal points, say, a_{1} and a_{2} with $a_{1} \leq a_{2}$, in the ε -level set of $\sigma_{\min}(F_{A}(r, \theta))$ with r varying and θ fixed, and then analogously, also obtain the two extremal level-set points b_{1} and b_{2} of $\sigma_{\min}(F_{B}(r, \theta))$ with $b_{1} \leq b_{2}$. By taking $r_{1} = \max\{a_{1}, b_{1}\}$ and $r_{2} = \min\{a_{2}, b_{2}\}$, we have that any global minimizer of v_{θ} must lie in $[r_{1}, r_{2}]$, since by construction, $v_{\theta}(r) > \varepsilon$ outside this interval. Thus, to obtain the value of $v(\theta)$, we simply solve two eigenvalue problems to obtain $[r_{1}, r_{2}]$ and then apply Chebfun to approximate v_{θ} on $[r_{1}, r_{2}]$ in order to obtain its globally minimal value.

Using Chebfun to approximate v over $[0, \pi)$, where for each θ , the value of $v(\theta)$ is also computed by applying Chebfun to v_{θ} , does lead to quite an expensive algorithm, as many evaluations of $\sigma_{\min}(F_A(r, \theta))$ and $\sigma_{\min}(F_B(r, \theta))$ for different values of θ and r are required. However, this nested Chebfun-based algorithm nevertheless has the virtue of being the very first algorithm to compute $\sup_{\lambda}^{V}(A, B)$, as opposed to just approximating it, e.g., within a factor of two by instead computing $\operatorname{sep}_{\lambda}^{D}(A, B)$.

Regarding the choice of z_0 , one might be tempted to use a local minimizer of f^V , but there are pros and cons to doing so. On the upside, if $\varepsilon = f^V(z_0)$ is close to $\sup_{\lambda}^V(A, B)$, v likely will be constant (with value ε) on a much of $[0, \pi)$, or all of it if $\varepsilon = \sup_{\lambda}^V(A, B)$, precisely because z_0 is a minimizer. This can greatly reduce the number of function evaluations required by Chebfun, but as discussed earlier in §5, functions with large constant portions can actually cause Chebfun to terminate prematurely. As such, we generally recommend that a minimizer of f^V not be used for z_0 .

Finally, recalling our recommendation at the end of Remark 6.7, when computing $\text{sep}_{\lambda}^{V}(A, B)$ via (8.4), we also similarly recommend refining Chebfun's result via subsequently applying local optimization. The upper bound given in (8.2) should also be computed.

9 Numerical experiments

All experiments were done in MATLAB R2021a on a computer with two Intel Xeon Gold 6130 processors (16 cores each, 32 total) and 192GB of RAM running CentOS Linux 7. We implemented our new methods using a recent build of Chebfun (commit 119f9ad) with splitting enabled and novectorcheck, and for simplicity, computed eigenvalues of C_{θ} and S_{θ} using eig in MATLAB; to account for rounding errors, the real part of any computed eigenvalue λ was set to zero if $|\text{Re }\lambda| \leq 10^{-8}$. For Algorithm 7.1, we used v1.6.4 of GRANSO with opt_tol=1e-14 to find local minimizers and used lines instead of rays for our globality certificates, as we observed that this was usually a bit faster. We forgo including any parallel processing experiments here, as we have previously validated the large benefits of using parallelism with our interpolation-based globality certificate approach in [Mit21, Section 5.2]. The codes used to generate the results in this paper are included in the supplementary materials, and we plan to add robust implementations to ROSTAPACK: RObust STAbility PACKage [Mitb].



Figure 9.1: For the example described in §9.1, pseudospectra of A(s) and B(s) (respectively solid and dotted contours) corresponding to Demmel's and Varah's versions of sep-lambda are shown along with the eigenvalues of A(s) and B(s) (respectively x's and dots) for $s \in \{10, 5, 0\}$. In the top right plot, $\Lambda_{\varepsilon}(A(s))$ and $\Lambda_{\varepsilon}(B(s))$ appear to touch at two places, but actually there is only one contact point (the one closer to the origin). In the three lower plots, Varah's sep-lambda is attained with $\varepsilon_1 = 0$.

9.1 An exploratory example

We first consider a simple example to explore the properties of our methods. We generated two different complex 10×10 matrices using **randn** and rescaled them so that the resulting matrices A and B both had spectral radii of 10. Then, for $s \in \{10, 5, 0\}$, we considered Demmel's and Varah's versions of sep-lambda for A(s) = A - sI and B(s) = B + sI. When s = 0, the spectra of A(s) and B(s) are "centered" are the origin, but when s is increased, the centers of the two spectra, -s and s, become more and more distant from each other; hence, on a macro level, increasing s generally increases the value of $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$, and this is always true once s becomes sufficiently large. Estimates of $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$ were computed using Algorithm 7.1, while estimates of $\operatorname{sep}_{\lambda}^{V}(A(s), B(s))$ were computed using both of our algorithms from §8; for Varah's sep-lambda, the estimates for both our algorithms agreed exactly since they were obtained at an eigenvalue of A(s). For Algorithm 7.1 and its extension to Varah's sep-lambda, we used $10 + 10\mathbf{i}$ as an initial point for optimization, which was chosen so that some restarts would be observed. In Fig. 9.1, we show the resulting pseudospectra of A(s) and B(s) at the perturbation levels given by $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$ and $\operatorname{sep}_{\lambda}^{V}(A(s), B(s))$.

We give performance statistics of Algorithm 7.1 on our exploratory example in Table 9.1. For s = 10, GRANSO found a global minimizer of $f^{\rm D}$ from the initial point and so only a single certificate computation was needed in this case, while two certificates were needed for the s = 5and s = 0 instances. On both of these, the first round of optimization only found a local minimizer, and so the first certificate instead returned new points to restart optimization. But as can be seen from Table 9.1, this happened with very little effort; only 15 evaluations of d_{ε} were needed to

Table 9.1: Performance data for Algorithm 7.1 for Demmel's version of sep-lambda on the example given in §9.1 for $s \in \{10, 5, 0\}$: " f^{D} evals." is the total number of evaluations of f^{D} during local optimization, "Certs." is the total number of certificates attempted, "All" and "Final" are the total number of evaluations of d_{ε} over all certificates and just the final one, respectively, and the final column is the total running time in seconds of Algorithm 7.1.

			d_{ε} e	vals.	Time (sec.)
s	$f^{\rm D}$ evals.	Certs.	All	Final	Alg. 7.1
10	156	1	2154	2154	4
5	150	2	8587	8572	5
0	205	2	22838	22823	13

Table 9.2: Performance data for our two algorithms described in §8 for Varah's version of seplambda on the example given in §9.1 for $s \in \{10, 5, 0\}$. For the extension of Algorithm 7.1 described in §8.1, " f^{V} evals." is the total number of evaluations of f^{V} during all optimization runs, "Certs." is the total number of certificates attempted, and " $d_{\varepsilon_1,\varepsilon_2}$ evals." is the total number of evaluations of $d_{\varepsilon_1,\varepsilon_2}$ over all certificates. For our algorithm described in §8.2, " f^{V} evals." is the total number of evaluations of f^{V} , "v evals." is the total number of evaluated in §8.2, " f^{V} evals." is the total number of evaluations of f^{V} , "v evals." is the total number of evaluated of the function v defined in (8.3). Finally, we report the total running time in seconds for each method respectively under the "§8.1" and "§8.1" columns.

	А	lg. from	§8.1	Alg. from	Time (sec.)		
s	$f^{\rm V}$ evals.	Certs.	$d_{\varepsilon_1,\varepsilon_2}$ evals.	$f^{\rm V}$ evals.	v evals.	§8.1	§8.2
10	188	1	4494	2512417	2044	2	139
5	142	1	4583	10161448	5900	2	548
0	150	1	6859	14736950	6502	2	777

find new starting points. Fig. 9.2 shows that the corresponding final configurations of d_{ε} are all nonnegative, as they should be when $\varepsilon = \sup_{\lambda}^{D}(A(s), B(s))$, per Corollary 6.6. Overall, we see that additional effort was needed to approximate d_{ε} as s is decreased, which is as we would expect because the behavior of d_{ε} generally becomes more complicated in proportion to how much the two (pseudo)spectra "intermingle", which for our test examples, is roughly controlled by s. Recalling that the search point z_0 defining d_{ε} is near the origin for these problems, this effect can be clearly observed by looking at Figs. 9.1a and 9.2 (and is also illustrated in Fig. 4.1, where each quadrant of the complex plane has a different amount of pseudospectral "intermingling"). For s = 10, the eigenvalues of A(s) and B(s) are separated from each other the most, which in turn leads to the final d_{ε} being rather straightforward; see Fig. 9.2a. However, the separation between the eigenvalues of A(s) and B(s) is reduced via making s smaller, and hence we see that d_{ε} becomes increasingly more complicated and with more discontinuities; see Fig. 9.2b and Fig. 9.2c.

Performance data for our two algorithms for Varah's sep-lambda are given in Table 9.2, where we see a similar effect with respect to changing shift s. However, the main takeaway here is that, as predicted, our algorithm from §8.2 is indeed many times slower than our extension of Algorithm 7.1 described in §8.1.

9.2 Comparing Algorithm 7.1 to the method of Gu and Overton

We now do a comparison of Algorithm 7.1 against the seplambda routine⁵, which is Overton's MATLAB implementation of his $\operatorname{sep}_{\lambda}^{D}(A, B)$ algorithm with Gu [GO06]. To do this, we generated two more examples in the manner as described in §9.1 but now for m = n = 20 and m = n = 40. For each, including our earlier m = n = 10 example, we computed $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$ for s = 0 and s = m = n using both our new method and seplambda. In order to obtain $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$

⁵Available at https://cs.nyu.edu/faculty/overton/software/seplambda/.



Figure 9.2: Each subfigure shows the final d_{ε} computed by Algorithm 7.1 for the example from §9.1 for $s \in \{10, 5, 0\}$. The components of d_{ε} are plotted as follows: $a_{\varepsilon} + b_{\varepsilon}$ (dotted) and d_{ε}^{AB} (solid); ℓ_{ε} does not appear as it is never negative when $\varepsilon = \operatorname{sep}_{\lambda}^{D}(A, B)$. The circle denotes the angle θ (with respect to z_{0}) associated with the best minimizer of f^{D} obtained and corresponds to the single place where $d_{\varepsilon}(\theta) = 0$, which is more easily seen in the \log_{10} plots on the right.

to high precision, we set the respective tolerances for both methods to 10^{-14} . For this comparison, we always initialized the first phase of optimization for our method from the origin.

A performance overview is reported in Table 9.3. In terms of accuracy, the estimates computed by our method for $\sup_{\lambda}^{D}(A(s), B(s))$ have high agreement with those computed by **seplambda**, though our method did return slightly better (lower) values for all the problems. On the non-

Table 9.3: Comparing Algorithm 7.1 and the method of Gu and Overton for Demmel's version of sep-lambda. The columns are the same as described in Table 9.1 except that we now additionally give the problem size under "m = n", the total running times in seconds of both methods, respectively "Alg. Algorithm 7.1" and "GO", and the relative difference between the estimates computed by both methods (Rel. Diff.), with positive indicating our method returned a better (lower) equal estimate for $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$.

				d_{ε} evals.		Time (sec.)		
m = n	s	$f^{\rm D}$ evals.	Certs.	All	Final	Alg. 7.1	GO	Rel. Diff.
10	10	65	1	2154	2154	3	16	$3.5 imes 10^{-14}$
10	0	75	1	23287	23287	14	15	0
20	20	97	1	4746	4746	15	1481	1.1×10^{-13}
20	0	346	3	31786	31756	78	1408	1.3×10^{-13}
40	40	128	2	5973	5910	95	237425	3.3×10^{-14}
40	0	295	3	29261	29231	407	230251	2.0×10^{-12}

shifted (s = 0) examples, our new method was 1.1 times faster than seplambda for m = n = 10, 18.1 times faster for m = n = 20, and 566.0 times faster for m = n = 40. Clearly, as the problems get larger, our method will be even faster relative to seplambda. For the shifted examples (s = m = n), the performance gaps are even wider: our new method was 6.2 times faster than seplambda for m = n = 10, 98.6 times faster for m = n = 20, and 2495.7 times faster for m = n = 40. The " d_{ε} evals." data for s = 0 and s = m = n in Table 9.3 for these problems also indicate that d_{ε} is generally less complex the more the eigenvalues of A and B are separated. Meanwhile, the running times of seplambda were relatively unchanged by the value of s, as shifting the eigenvalues of A and B has no direct effect on its computations. In Table 9.3, we can again infer that restarts in our method, when needed, happened with relatively few evaluations of d_{ε} . Per [Mit21, Section 5.2], the main work done in interpolation-based globality certificates is "embarrassingly parallel", and consequently, our method can further be accelerated by about an order of magnitude using parallel processing, and substantially more if minor tweaks are made to Chebfun to make it more amenable to parallelism.

Remark 9.1. Recalling the end of \S ² on possibly replacing the bisection phase of Gu and Overton's with optimization-with-restarts, we now empirically validate our claim that the benefit of such a modification is indeed quite limited and diminishes as the problem dimensions increase. Besides recording the total time to run seplambda on each problem for Table 9.3, we also recorded the time its initialization procedure required. Then, an upper bound for the best possible speedup is simply the total time divided by the initialization time, where we idealistically assume that opimizationwith-restarts has zero cost. For m = n respectively equal to 10, 20, and 40, the computed ratios were approximately 3.5, 2.2, and 1.6. Obviously, even these idealized speedups are nowhere near sufficient to overcome the very large performance gaps shown in Table 9.3 for m = n = 20, let alone m = n = 40, although such a modified version of seplambda would be close in performance to our method on the m = n = 10, s = 10 problem and likely pull ahead for the m = n = 10, s = 0problem. However, if we enabled parallel processing for Algorithm 7.1, then it would again be fastest on this problem too and probably by a large margin. Finally, note that if seplambda were further modified by also adapting the divide-and-conquer technique of $[GMO^+06]$, it still would be significantly slower than Algorithm 7.1, except for maybe the tiniest of problems. In the context of computing the distance to uncontrollability, we compared our interpolation-based globality certificates methodology with the method of $[GMO^+06]$, which uses both optimization-with-restarts and divide-and-conquer and also does not have any expensive initialization procedure, and our approach was roughly 5 to 43 times faster depending on the dimension; see [Mit21, Section 5.1].

9.3 Scaling performance of Algorithm 7.1

Finally, we examine the scaling performance of Algorithm 7.1 on some larger problems, which we constructed in the same fashion as before except that here we generated complex matrices A

Table 9.4: The columns are the same as described in Table 9.3 except that here we only give running times of Algorithm 7.1 and its computed estimates of $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$. The running time and accuracy comparisons with Gu and Overton's method are not provided since it would have taken far too long to run their method on these larger problems.

				$d_{arepsilon}$ e	vals.	Time (sec.)	
m = n	s	$f^{\rm D}$ evals.	Certs.	All	Final	Alg. 7.1	$\operatorname{sep}^{\mathrm{D}}_{\lambda}(A(s),B(s))$
100	100	89	1	5425	5425	704	$6.1677176880084 \times 10^{0}$
100	0	334	4	23689	23451	3045	$2.5004731832266 \times 10^{-2}$
200	200	210	2	23170	23155	11563	$6.3206868631252 \times 10^{0}$
200	0	319	2	20846	20831	10187	$4.2654521922541 \times 10^{-2}$
400	400	123	2	22456	22425	40966	$6.0394981396743 \times 10^{0}$
400	0	392	5	18237	18113	33931	$4.0258158186612\times 10^{-2}$
800	800	153	2	3113	3098	22112	$1.0584889222355 \times 10^{1}$
800	0	383	5	21962	19436	131887	$9.9483548512835 \times 10^{-3}$

and B via sprandn with a density of 0.1; this change was done solely to be able to store the matrices explicitly while keeping the file sizes small for up to m = n = 800. For these problem sizes, it was not feasible to attempt running Gu and Overton's method, so in Table 9.4, we only give performance data for Algorithm 7.1. The accuracy of each estimate ε for $\operatorname{sep}_{\lambda}^{D}(A(s), B(s))$ computed by Algorithm 7.1 was verified by creating a sufficiently high resolution plot of $\Lambda_{\varepsilon}(A(s))$ and $\Lambda_{\varepsilon}(B(s))$ and inspecting it to see whether or not the interiors of the two pseudospectra overlap. This visual check suffices to confirm the high accuracy of our new method because, per §7.4, local minimizers discovered on every iteration of Algorithm 7.1 will be computed to high accuracy, and the fact that $\varepsilon > \sup_{\lambda}^{D}(A(s), B(s))$ if and only if $\operatorname{int} \Lambda_{\varepsilon}(A(s)) \cap \operatorname{int} \Lambda_{\varepsilon}(B(s)) \neq \emptyset$; hence, to assess the accuracy of a computed estimate ε , we need only confirm whether or not Algorithm 7.1 converged to a global minimizer of $f^{\rm D}$ or only a local one, which is done by looking for the absence or presence, respectively, of pseudospectral overlap. For the pair of smallest problems (m = n = 100), Algorithm 7.1 respectively took about 11 and 50 minutes, while on the other extreme, Algorithm 7.1 needed about 6 and 37 hours, respectively, for the two m = n = 800problem instances. Again, using parallel processing can reduce these running times dramatically. Interestingly, for the intermediate sizes of m = n = 200 and m = n = 400, we actually see that Algorithm 7.1 was slightly more expensive on the instances with nonzero s, which suggests that the spectra of A(s) and B(s) for these particular examples would need to be shifted even further apart in order for the complexity of d_{ε} to decrease. Over all the problems tested, we see that Algorithm 7.1 required at most four restarts before converging, but once again, the costs of these restarts was generally negligible, with the one exception being the m = n = 800, s = 0 problem, where we can infer that the total cost of the four restarts was approximately 10% of the overall running time.

10 Concluding remarks

In this paper, we have introduced a new method to compute Demmel's version of sep-lambda that is much faster than the only previous known algorithm (due to Gu and Overton). Under our assumption that approximation of d_{ε} by interpolation is reliable, our method computes $\operatorname{sep}_{\lambda}^{D}(A, B)$ to arbitrary accuracy and generally behaves like a method with cubic work complexity, albeit one with a high constant factor. Nevertheless, our new approach is so much faster that it is now possible to calculate $\operatorname{sep}_{\lambda}^{D}(A, B)$ for moderately sized problems, e.g., for m, n in the thousands, which were simply intractable when using Gu and Overton's algorithm. We have also extended our algorithm to tackle Varah's version of sep-lambda. Although in this case global optimality cannot be guaranteed, the extension does rapidly compute locally optimal approximations which satisfy the necessary condition for global optimality. Furthermore, we have a proposed a second method to actually compute $\operatorname{sep}_{\lambda}^{V}(A, B)$, although this algorithm is significantly more expensive.

Acknowledgements

The author is grateful for the referees' helpful suggestions, and particularly for the first referee's observation that the efficient pseudospectral plotting technique of [Lui97] could be used in Algorithm 7.1 to bring down its worst-case asymptotic work complexity to cubic in n as opposed to quartic; see §7.2 and §7.6.

References

- [AGV17] G. Armentia, J.-M. Gracia, and F.-E. Velasco. On the boundaries of strict pseudospectra. Linear Algebra Appl., 525:20–32, 2017.
- [BB90] S. Boyd and V. Balakrishnan. A regularity result for the singular values of a transfer matrix and a quadratically convergent algorithm for computing its L_{∞} -norm. Systems Control Lett., 15(1):1–7, 1990.
- [BBK89] S. Boyd, V. Balakrishnan, and P. Kabamba. A bisection method for computing the \mathcal{H}_{∞} norm of a transfer matrix and related problems. *Math. Control Signals Systems*, 2:207–219, 1989.
- [BCL⁺20] J. V. Burke, F. E. Curtis, A. S. Lewis, M. L. Overton, and L. E. A. Simões. Gradient sampling methods for nonsmooth optimization. In A. M. Bagirov, M. Gaudioso, N. Karmitsa, M. M. Mäkelä, and S. Taheri, editors, *Numerical Nonsmooth Optimization: State of the Art Algorithms*, pages 201–225, Cham, 2020. Springer International Publishing.
- [BLO03] J. V. Burke, A. S. Lewis, and M. L. Overton. Robust stability and a criss-cross algorithm for pseudospectra. IMA J. Numer. Anal., 23(3):359–375, 2003.
- [BLO04] J. V. Burke, A. S. Lewis, and M. L. Overton. Pseudospectral components and the distance to uncontrollability. SIAM J. Matrix Anal. Appl., 26(2):350–361, 2004.
- [BL005] J. V. Burke, A. S. Lewis, and M. L. Overton. A robust gradient sampling algorithm for nonsmooth, nonconvex optimization. SIAM J. Optim., 15(3):751–779, 2005.
- [BM18] P. Benner and T. Mitchell. Faster and more accurate computation of the \mathcal{H}_{∞} norm via optimization. SIAM J. Sci. Comput., 40(5):A3609–A3635, October 2018.
- [BM19] P. Benner and T. Mitchell. Extended and improved criss-cross algorithms for computing the spectral value set abscissa and radius. *SIAM J. Matrix Anal. Appl.*, 40(4):1325–1352, 2019.
- [BMX99] P. Benner, V. Mehrmann, and H. Xu. A note on the numerical solution of complex Hamiltonian and skew-Hamiltonian eigenvalue problems. *Electron. Trans. Numer. Anal.*, 8:115–126, 1999.
- [BS90] N. A. Bruinsma and M. Steinbuch. A fast algorithm to compute the H_{∞} -norm of a transfer function matrix. Systems Control Lett., 14(4):287–293, 1990.
- [BSV12] P. Benner, V. Sima, and M. Voigt. \mathcal{L}_{∞} -norm computation for continuous-time descriptor systems using structured matrix pencils. *IEEE Trans. Autom. Control*, 57(1):233–238, January 2012.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.
- [Bye88] R. Byers. A bisection method for measuring the distance of a stable matrix to unstable matrices. *SIAM J. Sci. Statist. Comput.*, 9:875–881, 1988.
- [Bye90] R. Byers. Detecting nearly uncontrollable pairs. In M. A. Kaashoek, J. H. Schuppen, and A. C. M. Ran, editors, Signal processing, scattering and operator theory, and numerical methods, volume 3 of Proc. of the International Symposium MTNS-89, Amsterdam 1989, pages 447-457, Boston, MA, 1990. Birkhäuser.
- [CMO17] F. E. Curtis, T. Mitchell, and M. L. Overton. A BFGS-SQP method for nonsmooth, nonconvex, constrained optimization and its evaluation using relative minimization profiles. Optim. Methods Softw., 32(1):148–181, 2017.
- [Dem83] J. Demmel. A Numerical Analyst's Jordan Canonical Form. PhD thesis, University of California, Berkeley, Berkeley, CA 94720, 1983.
- [Dem86] J. Demmel. Computing stable eigendecompositions of matrices. Linear Algebra Appl., 79:163– 193, 1986.
- [Dem87] J. Demmel. A counterexample for two conjectures about stability. IEEE Trans. Autom. Control, 32(4):340–342, April 1987.

- [DHT14] T. A Driscoll, N. Hale, and L. N. Trefethen. Chebfun Guide. Pafnuty Publications, Oxford, UK, 2014.
- [GMO⁺06] M. Gu, E. Mengi, M. L. Overton, J. Xia, and J. Zhu. Fast methods for estimating the distance to uncontrollability. SIAM J. Matrix Anal. Appl., 28(2):477–502, 2006.
- [GN93] M. Gao and M. Neumann. A global minimum search algorithm for estimating the distance to uncontrollability. *Linear Algebra Appl.*, 188/189:305–350, 1993.
- $[GO06] M. Gu and M. L. Overton. An algorithm to compute Sep_{<math>\lambda$}. SIAM J. Matrix Anal. Appl., 28(2):348–359, 2006.
- [Gu00] M. Gu. New methods for estimating the distance to uncontrollability. SIAM J. Matrix Anal. Appl., 21(3):989–1003, 2000.
- [Hig93] N. J. Higham. Perturbation theory and backward error for AX XB = C. BIT, 33(1):124–136, 1993.
- [HW97] C. He and G. A. Watson. An algorithm for computing the numerical radius. IMA J. Numer. Anal., 17(3):329–342, June 1997.
- [KK14] M. Karow and D. Kressner. On a perturbation bound for invariant subspaces of matrices. SIAM J. Matrix Anal. Appl., 35(2):599–618, 2014.
- [LO13] A. S. Lewis and M. L. Overton. Nonsmooth optimization via quasi-Newton methods. Math. Program., 141(1–2, Ser. A):135–163, 2013.
- [Lui97] S. H. Lui. Computation of pseudospectra by continuation. SIAM J. Sci. Comput., 18(2):565– 573, 1997.
- [Mita] T. Mitchell. GRANSO: GRadient-based Algorithm for Non-Smooth Optimization. http://timmitchell.com/software/GRANSO. See also [CMO17].
- [Mitb] T. Mitchell. ROSTAPACK: RObust STAbility PACKage. http://timmitchell.com/software/ROSTAPACK.
- [Mit20] T. Mitchell. Computing the Kreiss constant of a matrix. SIAM J. Matrix Anal. Appl., 41(4):1944–1975, 2020.
- [Mit21] T. Mitchell. Fast interpolation-based globality certificates for computing Kreiss constants and the distance to uncontrollability. *SIAM J. Matrix Anal. Appl.*, 42(2):578–607, 2021.
- [Mit22] T. Mitchell. Convergence rate analysis and improved iterations for numerical radius computation. *To appear in SIAM J. Sci. Comput.*, Accepted October, 2022. Preprint available at https://arxiv.org/abs/2002.00080.
- [MO05] E. Mengi and M. L. Overton. Algorithms for the computation of the pseudospectral radius and the numerical radius of a matrix. *IMA J. Numer. Anal.*, 25(4):648–669, 2005.
- [MVD20a] V. Mehrmann and P. Van Dooren. Optimal robustness of passive discrete-time systems. IMA J. Math. Control Inform., 37(4):1248–1269, July 2020.
- [MVD20b] V. Mehrmann and P. M. Van Dooren. Optimal robustness of port-Hamiltonian systems. SIAM J. Matrix Anal. Appl., 41(1):134–151, 2020.
- [MVD22] T. Mitchell and P. Van Dooren. Root-max problems, hybrid expansion-contraction, and quadratically convergent optimization of passive systems. To appear in SIAM J. Matrix Anal. Appl., Accepted November, 2022. Preprint available at http://arxiv.org/abs/2109.00974.
- [PPT09] R. Pachón, R. B. Platte, and L. N. Trefethen. Piecewise-smooth chebfuns. IMA J. Numer. Anal., 30(4):898–916, July 2009.
- [RKBA21] N. Roy, M. Karow, S. Bora, and G. Armentia. Approximation of pseudospectra of block triangular matrices. *Linear Algebra Appl.*, 623:398–419, 2021. Special issue in honor of Paul Van Dooren.
- [Sim96] V. Simoncini. On the numerical solution of AX XB = C. BIT, 36(4):814–830, 1996.
- [Ste73] G. W. Stewart. Error and perturbation bounds for subspaces associated with certain eigenvalue problems. SIAM Rev., 15:727–764, 1973.
- [TE05] L. N. Trefethen and M. Embree. Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators. Princeton University Press, Princeton, 2005.
- [Tre20] L. N. Trefethen. Personal communication, August 2020. Recorded at https://groups.google.com/g/chebfun-users/c/9d6zzPkfwZk/m/cJJKx3C4AQAJ.
- [Var67] J. M. Varah. The Computation of Bounds for the Invariant Subspaces of a General Matrix Operator. PhD thesis, Stanford University, Palo Alto, CA 94305, 1967.
- [Var79] J. M. Varah. On the separation of two matrices. SIAM J. Numer. Anal., 16(2):216–222, 1979.