

Subspace acceleration for the Crawford number and related eigenvalue optimization problems

Daniel Kressner^{*}, Ding Lu[†] and Bart Vandereycken[‡]

April 25, 2017

Abstract

This paper is concerned with subspace acceleration techniques for computing the Crawford number, that is, the distance between zero and the numerical range of a matrix A . Our approach is based on an eigenvalue optimization characterization of the Crawford number. We establish local convergence of order $1 + \sqrt{2} \approx 2.4$ for an existing subspace method applied to such and other eigenvalue optimization problems involving a Hermitian matrix that depends analytically on one parameter. For the particular case of the Crawford number, we show that the relevant part of the objective function is strongly concave. In turn, this enables us to develop a subspace method that only uses three-dimensional subspaces but still achieves global convergence and a local convergence that is at least quadratic. A number of numerical experiments confirm our theoretical results and reveal that the established convergence orders appear to be tight.

1 Introduction

During the last few years, subspace methods for optimizing eigenvalues of Hermitian matrices or singular values of non-Hermitian matrices have been developed, analyzed, and applied to a variety of problems. A common trait of these methods is that they use the subspace spanned by eigenvectors or singular vectors from previous iterations to build a low-dimensional surrogate model that determines the next iterate. Examples include subspace methods for computing (real) pseudospectral abscissa for linear [15, 18] and nonlinear [19] eigenvalue problems, for computing the \mathcal{H}_∞ norm of a linear control system [1], as well as for solving Hermitian eigenvalue optimization problems depending on several parameters [13].

In this work, we revisit eigenvalue optimization problems depending on one parameter, with a particular focus on their application to Crawford number computation. The Crawford number [5, 24] of a matrix $A \in \mathbb{C}^{n \times n}$ is defined as the distance of its numerical range $\mathcal{F}(A)$ from zero:

$$\gamma(A) = \min\{|z| : z \in \mathcal{F}(A)\}, \quad \mathcal{F}(A) = \{v^*Av : v \in \mathbb{C}^n, \|v\|_2 = 1\}. \quad (1)$$

This quantity plays an important role in the sensitivity analysis of matrix definite pairs [4, 24], and the detection of hyperbolic quadratic eigenvalue problems [7, 11]. Possibly even more importantly, $\gamma(A)$ and its generalization to bilinear forms on Hilbert spaces plays the role of the coercivity

^{*}EPFL, Section of Mathematics, Lausanne, Switzerland. (daniel.kressner@epfl.ch)

[†]Department of Mathematics, University of Geneva, Geneva, Switzerland. (Ding.Lu@unige.ch)

[‡]Department of Mathematics, University of Geneva, Geneva, Switzerland. (Bart.Vandereycken@unige.ch)

constant for (discretized) non-self-adjoint differential/integral operators and their discretization; see [2] for a recent example.

By the Hausdorff–Toeplitz theorem, the numerical range $\mathcal{F}(A)$ is a compact convex set in \mathbb{C} . Hence, (1) is a convex optimization problem and a number of algorithms explicitly exploit this property, for example, by approximating the numerical range as a sequence of convex polygons [26]. However, it turns out to be more effective to avoid such geometric considerations in $\mathbb{C} \cong \mathbb{R} \times \mathbb{R}$ and exploit that (1) is in fact a univariate optimization problem. To see this, we first write

$$A = S + \iota K,$$

where $S = (A + A^*)/2$ and $K = (A - A^*)/2\iota$ are both Hermitian matrices. Using the variational characterization of eigenvalues of Hermitian matrices, one then obtains

$$\gamma(A) = \max \left\{ \max_{\theta \in [0, 2\pi]} \lambda_{\min}(S \cos \theta + K \sin \theta), \quad 0 \right\}, \quad (2)$$

see, e.g., [4, Thm. 2.1] and [11, eq. (2.8)]. Based on this formula, a level set algorithm [4, 11] and bisection algorithms [3, 8] were developed. While being reliable, these algorithms are only linearly convergent and they typically require a relatively large number of iterations to attain high accuracy.

Contributions. The approach considered in this paper consists of applying existing subspace methods, specifically the one from [13], to (2) and, more generally, to univariate eigenvalue optimization problems. We would like to emphasize the following two novel contributions:

Existing convergence analysis [13, 15, 18] of subspace methods for eigenvalue or singular value optimization establish local superlinear convergence. The restriction to the univariate case enables the use of univariate analytic interpolation results, which in turn allows us to derive in Section 2 a much stronger result: local convergence of order $\sigma = 1 + \sqrt{2} \approx 2.4$. In contrast to [13], we also do not need to resort to an “extended” and more expensive version of the subspace method in order to establish this result. Numerical experiments reveal that our obtained value for σ appears to be tight.

Restricting ourselves further to the specific eigenvalue optimization problem (2), we establish in Section 3 a concavity property for the objective function and improve upon a result from [8]. This allows us to show that subspace methods enjoy a rather useful “bracketing” property when applied to Crawford number computations. Based upon these insights, we develop a “3-vector” subspace method that works with only three-dimensional subspaces but still achieves global convergence and a local convergence that is quadratic or even higher.

2 Hermitian eigenvalue optimization problem

In the following, we consider the univariate eigenvalue optimization problem

$$\max_{\theta \in \Omega} \varphi(\theta) \quad \text{with} \quad \varphi(\theta) = \lambda_{\min}(H(\theta)), \quad (3)$$

where $\Omega \subset \mathbb{R}$ is a closed interval, and $H(\theta) \in \mathbb{C}^{n \times n}$ is Hermitian and real analytic for $\theta \in \Omega$, that is, it admits an analytic extension in an open neighborhood $\Omega_{\mathbb{C}} \subset \mathbb{C}$.

Given an orthonormal basis $V \in \mathbb{C}^{n \times k}$ of a k -dimensional subspace of \mathbb{C}^n , the subspace acceleration techniques discussed, e.g., in [13, 22], proceed by considering the reduced objective function

$$\varphi(\theta; V) = \lambda_{\min}(V^* H(\theta) V).$$

Typically $k \ll n$ and it can therefore be expected that computing the maximum of $\varphi(\theta; V)$ on Ω is cheaper than solving (3). In particular, this is the case when H admits an affine linear decomposition $H(\theta) = \sum_{\ell=1}^m f_\ell(\theta)H_\ell$ with scalar functions f_ℓ and Hermitian matrices H_ℓ such that $m \ll n$. After precomputing $V^*H_\ell V$ for $\ell = 1, \dots, m$, each evaluation of $\varphi(\theta; V)$ only requires the addition of m matrices and the solution of one $k \times k$ eigenvalue problem.

The following Lemma 1 recalls results from [13, 22] on basic properties of the reduced function $\varphi(\theta; V)$; its proof is included for completeness.

Lemma 1. *Let U and V be orthonormal bases of subspaces $\mathcal{U} \subset \mathcal{V}$ of \mathbb{C}^n . Then the following properties hold.*

(a) *Monotonicity:* $\varphi(\theta; U) \geq \varphi(\theta; V) \geq \varphi(\theta)$.

(b) *Interpolation:* Let $\theta \in \Omega$ and suppose that \mathcal{V} contains an eigenvector belonging to the eigenvalue $\lambda_{\min}(H(\theta))$ of $H(\theta)$. Then, $\varphi(\theta; V) = \varphi(\theta)$. In addition, if $\lambda_{\min}(H(\theta))$ is a simple eigenvalue then $\varphi'(\theta; V) = \varphi'(\theta)$, where all derivatives are taken with respect to θ .

Proof. (a) Monotonicity follows directly from the variational characterization of eigenvalues:

$$\varphi(\theta; U) = \min_{u \in \mathcal{U}} \frac{u^* H(\theta) u}{u^* u} \geq \min_{u \in \mathcal{V}} \frac{u^* H(\theta) u}{u^* u} = \varphi(\theta; V). \quad (4)$$

The second inequality is obtained from the first since $\varphi(\theta) = \varphi(\theta; I_n)$.

(b) Let $v(\theta)$ be an eigenvector belonging to $\lambda_{\min}(H(\theta))$. Since $v(\theta) \in \mathcal{V} \setminus \{0\}$ there is $x \in \mathbb{C}^k$ such that $v(\theta) = Vx$. Then, by definition of the eigenvector $v(\theta)$, we have

$$\varphi(\theta; V) = \min_{z \in \mathbb{C}^k} \frac{z^* V^* H(\theta) V z}{z^* z} \leq \frac{x^* V^* H(\theta) V x}{x^* x} = \frac{v(\theta)^* H(\theta) v(\theta)}{v(\theta)^* v(\theta)} = \varphi(\theta).$$

Combined with (a), we obtain $\varphi(\theta; V) = \varphi(\theta)$. If $\varphi(\theta)$ is a simple eigenvalue of $H(\theta)$, the Cauchy interlacing theorem yields that $\varphi(\theta; V)$ is also a simple eigenvalue of $V^* H(\theta) V$, with eigenvector x . By well-known results on the derivatives of eigenvalues (see, e.g., [16]), we obtain

$$\varphi'(\theta) = \frac{v(\theta)^* H'(\theta) v(\theta)}{v(\theta)^* v(\theta)} = \frac{x^* V^* H'(\theta) V x}{x^* x} = \varphi'(\theta; V). \quad \square$$

By Lemma 1, $\varphi(\theta; V)$ provides an upper bound to $\varphi(\theta)$ that monotonically decreases as the dimension of \mathcal{V} increases. In view of the intimate relation between gradients of eigenvalues and eigenvectors, it is a natural idea to extend V by an eigenvector corresponding to $\varphi(\theta; V)$. This idea leads to the basic subspace method from [13], summarized in Algorithm 1. The difficulty of solving the reduced problem in line 4 depends on the application. It turns out to be quite cheap for the Crawford number, as we will see in Section 3. Other applications with this property are also listed in [13].

In the exceptional situation that Algorithm 1 stagnates it has found a global maximum.

Lemma 2. *If $\theta_{k+1} = \theta_j$ for some $j \leq k$ then Algorithm 1 terminates and $\theta_{k+1}, \varphi_{k+1}$ solve (3).*

Proof. Without loss of generality, we may assume $j = k$. By construction, λ_k is a lower bound and φ_{k+1} is an upper bound for $\max_{\theta \in \Omega} \varphi(\theta)$. Now, if $\theta_k = \theta_{k+1}$ then

$$\varphi_{k+1} = \lambda_{\min}(V_k^* H(\theta_{k+1}) V_k) = \lambda_{\min}(V_k^* H(\theta_k) V_k) = \lambda_{\min}(H(\theta_k)) = \lambda_k,$$

where we used $v_k \in \text{span}(V_k)$ in the second-to-last equality. □

Algorithm 1 Subspace method for univariate eigenvalue optimization

Input: Real analytic Hermitian matrix-valued function $H(\theta)$, initial guess θ_0 , tolerance $\text{tol} > 0$.

Output: Approximation solution $\theta_{k+1}, \varphi_{k+1}$ of eigenvalue optimization problem (3).

- 1: Compute $\lambda_0 = \lambda_{\min}(H(\theta_0))$ and corresponding normalized eigenvector v_0 .
 - 2: Initialize $V_0 = v_0$.
 - 3: **for** $k = 0, 1, \dots, n - 1$ **do**
 - 4: Solve $\theta_{k+1} = \arg \max_{\theta \in \Omega} \lambda_{\min}(V_k^* H(\theta) V_k)$ and set $\varphi_{k+1} = \lambda_{\min}(V_k^* H(\theta_{k+1}) V_k)$.
 - 5: *Stopping criteria:* **if** $\varphi_{k+1} - \lambda_k < \text{tol} \cdot |\varphi_{k+1}|$ **then** terminate.
 - 6: Compute the smallest eigenvalue λ_{k+1} with normalized eigenvector v_{k+1} of $H(\theta_{k+1})$.
 - 7: *Subspace update:* $V_{k+1} = \text{orth}([V_k, v_{k+1}])$.
 - 8: **end for**
-

A consequence of this lemma is that if $v_{k+1} \in \text{span}(V_k)$, we have $V_{k+1} = V_k$ and $\theta_{k+1} = \theta_k$. This implies we have solved (3). Otherwise, Algorithm 1 converges trivially after at most n steps, since V_{n-1} will be a basis of \mathbb{C}^n and Lemma 1 shows that φ_n solves (3) exactly. In [13, Theorem 3.1], an infinite-dimensional setting was considered to meaningfully establish a global convergence result, in the sense that every accumulation point of the sequence $(\theta_k)_{k=0}^\infty$ is in fact a global maximum of the optimization problem (3).

2.1 Local convergence

Here and in the following we assume that the eigenvalues corresponding to $\varphi(\theta_i)$ in Algorithm 1 are simple and hence differentiable. By the interpolation property established in Lemma 1(b), the reduced function $\varphi(\theta; V_k)$ from the k th iteration of Algorithm 1 then constitutes a (non-polynomial) Hermite interpolation of $\varphi(\theta)$, that is, it interpolates both function values and derivatives at the sampling points:

$$\varphi(\theta_i; V_k) = \varphi(\theta_i) \quad \text{and} \quad \varphi'(\theta_i; V_k) = \varphi'(\theta_i), \quad \text{for } i = 0, 1, \dots, k. \quad (5)$$

From (5) we can therefore expect a highly accurate approximation when the iterates θ_i are close to a global maximum θ_* of (3). Moreover, this process has a snowball effect: the accuracy of all previous $\theta_0, \dots, \theta_k$ accumulates to produce the next, even more accurate iterate θ_{k+1} . The following theorem quantifies this effect. Its proof is deferred to Section 2.2.

Theorem 1 (Error recurrence). *Let θ_* solve (3) and assume that $\varphi(\theta_*) = \lambda_{\min}(H(\theta_*))$ is a simple eigenvalue and $\varphi''(\theta_*) < 0$. Suppose that Algorithm 1 starts with $\theta_0 \in [\theta_* - r, \theta_* + r]$ for sufficiently small $r > 0$, and all subsequent iterates $\theta_1, \theta_2, \theta_3, \dots$ remain in $[\theta_* - r, \theta_* + r]$. Then*

$$|\theta_{k+1} - \theta_*| \leq |\theta_k - \theta_*| \quad \text{for all } k \geq 1. \quad (6)$$

Moreover, there exists a constant $R > 0$ independent of k and θ_0 such that the weighted error $e_\ell = |\theta_\ell - \theta_*|/R$ satisfies

$$e_{k+1} \leq C \cdot (k+1) \cdot e_k \prod_{\ell=0}^{k-1} e_\ell^2, \quad k = 1, 2, \dots, \quad (7)$$

where $C = 16M/|\varphi''(\theta_*)R|$ with $M = \max_{z \in \mathbb{C}: |z-\theta_*|=R} \|H(z)\|_2$.

Remark 1. Note that the local convergence analysis in [13] in the context of multivariate eigenvalue optimization is carried out for an extended subspace algorithm, a variant of Algorithm 1 that uses additional vectors (besides an eigenvector of $H(\theta_k)$) to expand the subspace V_k . It was shown in [13] that this extended subspace algorithm converges locally superlinearly. In our setting and analysis, these additional vectors are unnecessary and, moreover, we prove a stronger result than superlinear convergence.

The following result shows that the recurrence (7) from Theorem 1 corresponds to high-order convergence of the error to zero.

Theorem 2. Consider a sequence $(e_k)_{k=0}^\infty$ satisfying the recurrence (7), with $3Ce_0^2 < 1$ and $2e_1 \leq 1$. Then the sequence converges to zero with a local R -convergence order

$$\sigma = 1 + \sqrt{2} \approx 2.4142. \quad (8)$$

In other words, there is a sequence $(\varepsilon_k)_{k=0}^\infty$ such that $e_k \leq \varepsilon_k$ and $\lim_{k \rightarrow \infty} |\varepsilon_k|/|\varepsilon_{k-1}|^\sigma < \infty$.

Proof. Introducing the auxiliary terms $\tilde{e}_\ell = (\ell + 1)e_\ell$, we get from the error recurrence (7) that

$$\tilde{e}_{k+1} \leq 3C \cdot \tilde{e}_k \prod_{\ell=0}^{k-1} \tilde{e}_\ell^2, \quad k = 1, 2, \dots$$

Taking the logarithm on both sides and replacing the inequality by an equality we obtain that $\tilde{e}_k \leq \varepsilon_k$, with the upper bound determined by $\varepsilon_0 = \tilde{e}_0$, $\varepsilon_1 = \tilde{e}_1$ and the recurrence

$$\ln \varepsilon_{k+1} = \ln \varepsilon_k + 2 \sum_{\ell=0}^{k-1} \ln \varepsilon_\ell + \ln 3C, \quad k = 1, 2, \dots \quad (9)$$

Setting $s_k = \frac{1}{2} \ln 3C + \sum_{\ell=0}^k \ln \varepsilon_\ell$, this becomes equivalent to

$$s_{k+1} = 2s_k + s_{k-1}, \quad s_0 = \ln e_0 + \frac{1}{2} \ln 3C < 0, \quad s_1 = s_0 + \ln 2e_1 < 0.$$

By standard techniques for solving difference equations, there exist c_1, c_2 such that

$$s_k = c_1(1 + \sqrt{2})^k + c_2(1 - \sqrt{2})^k.$$

Observe that $s_k < 0$ for all k . Since the first term dominates when $k \rightarrow \infty$, we have $c_1 < 0$. From

$$\frac{\ln \varepsilon_k}{\sigma^k} = \frac{s_k - s_{k-1}}{\sigma^k} \xrightarrow{k \rightarrow \infty} \frac{c_1 \sqrt{2}}{\sigma} < 0,$$

one obtains Q-convergence of ε_k with order σ . Since $e_k \leq \tilde{e}_k \leq \varepsilon_k$ for all k , this shows the claimed R-convergence of e_k . \square

2.2 Proof of Theorem 1

Throughout this section, we suppose that the assumptions of Theorem 1 hold. In particular, $\varphi(\theta_*) = \lambda_{\min}(H(\theta_*))$ is a simple eigenvalue of $H(\theta_*)$ and $\varphi''(\theta_*) < 0$ holds. In addition, $H(z)$ is complex analytic on $\Omega_{\mathbb{C}}$.

In the following, we measure the distance between a vector $u \in \mathbb{C}^n$ and the subspace spanned by $V \in \mathbb{C}^{n \times k}$ as

$$\text{dist}(V, u) = \min \{ \|u - v\|_2 : v \in \text{span}(V) \}. \quad (10)$$

We let

$$v(\theta) = \text{a normalized eigenvector belonging to } \lambda_{\min}(H(\theta)). \quad (11)$$

The lemma below shows that the reduced function $\varphi(\theta; V)$ can be analytically extended to a region around θ_* . Moreover, this region can be chosen uniformly over all subspaces V containing a sufficiently good approximation of $v(\theta_*)$.

Lemma 3 (Analyticity). *Let V be an orthonormal basis satisfying $\text{dist}(V, v(\theta_*)) \leq \delta$. Provided that $\delta > 0$ is sufficiently small, there exists a constant $R > 0$ (independent of V) such that the following statements hold:*

(a) *The function*

$$\widehat{\varphi}(z; V) = \lambda_*(V^*H(z)V), \quad (12)$$

where λ_ denotes the eigenvalue closest to $\varphi(\theta_*; V)$, is well defined on the disc $D_{\mathbb{C}}(\theta_*; R) = \{z \in \mathbb{C} : |z - \theta_*| \leq R\}$;*

(b) *$\widehat{\varphi}(z; V)$ is complex analytic on $D_{\mathbb{C}}(\theta_*; R)$;*

(c) *$\widehat{\varphi}(\theta; V) = \varphi(\theta; V)$ for $\theta \in \mathbb{R} \cap D_{\mathbb{C}}(\theta_*; R)$. Hence, $\widehat{\varphi}(\cdot; V)$ is the analytic continuation of $\varphi(\cdot; V)$.*

Proof. (a) and (b) Let $\lambda_{\min-1}(\cdot)$ denote the second smallest eigenvalue of a Hermitian matrix. By our assumptions,

$$\text{gap} = \lambda_{\min-1}(H(\theta_*)) - \lambda_{\min}(H(\theta_*)) > 0.$$

We claim this induces a nonzero gap for the reduced matrix $V^*H(\theta_*)V$ provided that δ is sufficiently small. To see this, let $v(\theta_*) = v + d$ with $v \in \text{span}(V)$ and $\|d\|_2 = \text{dist}(V, v(\theta_*)) \leq \delta$. Clearly, v depends continuously on δ around zero and hence $v^*H(\theta_*)v/v^*v \rightarrow v(\theta_*)^*H(\theta_*)v(\theta_*) = \lambda_1(H(\theta_*))$ as $\delta \rightarrow 0$. We can therefore choose $\delta > 0$ such that

$$\lambda_{\min}(V^*H(\theta_*)V) = \min_{x \in \mathcal{V}} \frac{x^*H(\theta_*)x}{x^*x} \leq \frac{v^*H(\theta_*)v}{v^*v} < \lambda_{\min}(H(\theta_*)) + \text{gap}/2.$$

In turn, using eigenvalue interlacing, the gap for the reduced problem satisfies

$$\begin{aligned} \widetilde{\text{gap}} &= \lambda_{\min-1}(V^*H(\theta_*)V) - \lambda_{\min}(V^*H(\theta_*)V) \\ &> \lambda_{\min-1}(H(\theta_*)) - \lambda_{\min}(H(\theta_*)) - \text{gap}/2 = \text{gap}/2. \end{aligned} \quad (13)$$

In particular, $\lambda_{\min}(V^*H(\theta_*)V)$ is a simple eigenvalue.

We now choose $R > 0$ sufficiently small such that $D_{\mathbb{C}}(\theta_*; R)$ is contained in the domain of analyticity $\Omega_{\mathbb{C}}$ of $H(z)$ and, moreover,

$$\|H(z) - H(\theta_*)\|_2 \leq \text{gap}/4, \quad \forall z \in D_{\mathbb{C}}(\theta_*; R).$$

Because V is an orthonormal basis, this also implies

$$\|V^*H(z)V - V^*H(\theta_*)V\|_2 \leq \text{gap}/4, \quad \forall z \in D_{\mathbb{C}}(\theta_*; R). \quad (14)$$

Note that $V^*H(\theta_*)V$ is Hermitian while $V^*H(z)V$ is not. Applying classical eigenvalue perturbation results for *nearly Hermitian* matrices (see e.g. [12] and [25, Thm 5.1, and the discussion below the proof]) implies that there is exactly one eigenvalue of $V^*H(z)V$ in the disc $D_{\mathbb{C}}(\lambda_{\min}(V^*H(\theta_*)V); \text{gap}/4)$ for each $z \in D_{\mathbb{C}}(\theta_*; R)$, thus showing (a), that the eigenvalue function $\lambda_*(V^*H(z)V)$ is well defined on $D_{\mathbb{C}}(\theta_*; R)$. This function is complex analytic due to the analyticity of simple eigenvalues [14]. Since the choice of R to satisfy (14) is independent of V , once δ is sufficiently small, we have thus proved (b).

(c) The result follows immediately from the continuity of $\lambda_{\min}(V^*H(\theta)V)$ for $\theta \in \mathbb{R}$, and the uniqueness of the eigenvalue $\lambda_*(V^*H(\theta)V)$ in the interval $\theta \in [\theta_* - R, \theta_* + R] \subset D_{\mathbb{C}}(\theta_*, R)$. \square

Trivially, because of $\varphi(\theta) = \varphi(\theta; I_n)$, Lemma 3 implies that $\varphi(\theta)$ also has an analytic extension to $\widehat{\varphi}(z) = \widehat{\varphi}(z; I_n)$ in the complex disc $D_{\mathbb{C}}(\theta_*; R)$ with the same estimate for R .

Lemma 4 (Analytic approximation). *Let δ and R be as in Lemma 3. Then there exists $0 < r \leq R/2$ such that the following statements holds.*

(i) $\text{dist}(v(\theta_*), v(\theta)) \leq \delta$ for all $\theta \in \mathbb{R}$ with $|\theta - \theta_*| \leq r$.

(ii) Consider any $k + 1$ mutually distinct points $\{\theta_i\}_{i=0}^k \subset [\theta_* - r, \theta_* + r]$ ordered in decreasing distance to θ_* , and suppose that $v(\theta_0), v(\theta_1), \dots, v(\theta_k) \in \text{span}(V)$. Then:

(a) Both $\varphi(\theta)$ and $\varphi(\theta; V)$ are real analytic on $[\theta_* - R, \theta_* + R]$.

(b) With $C = 16M/R$ and $M = \max_{z \in \mathbb{C}: |z - \theta_*| = R} \|H(z)\|_2$,

$$|\varphi'(\theta_*; V) - \varphi'(\theta_*)| \leq C \cdot (k + 1) \cdot (R/2)^{-(2k+1)} \cdot |\theta_* - \theta_k| \prod_{i=0}^{k-1} |\theta_* - \theta_i|^2. \quad (15)$$

(c) For all $\theta \in \mathbb{R}$ with $|\theta - \theta_*| \leq r$,

$$\varphi''(\theta; V) \leq \varphi''(\theta_*)/2 < 0. \quad (16)$$

Proof. (i) This follows immediately from the fact that the eigenvalue $\lambda_{\min}(H(\theta_*))$ is simple, which implies the eigenvector $v(\theta)$ can be chosen to be continuous around θ_* . Hence, $v(\theta) \rightarrow v(\theta_*)$ as $\theta \rightarrow \theta_*$; see, e.g., [14].

(iia) This follows from Lemma 3, since $\text{dist}(V, v(\theta_*)) \leq \text{dist}(v(\theta_i), v(\theta_*)) \leq \delta$ according to (i).

(iib) We first recall from Lemma 1 that

$$\varphi(\theta_i; V) = \varphi(\theta_i), \quad \varphi'(\theta_i; V) = \varphi'(\theta_i), \quad \text{for } i = 0, \dots, k.$$

Let us reorder the sequence $\{\theta_k\}_{i=0}^k$ such that $\theta_{\Pi(0)} < \theta_{\Pi(1)} < \dots < \theta_{\Pi(k)}$, where Π is an appropriately chosen permutation. Then the first interpolation condition combined with Rolle's theorem implies that there exist $\widehat{\theta}_i$ with

$$\theta_{\Pi(i)} < \widehat{\theta}_i < \theta_{\Pi(i+1)} \quad \text{and} \quad \varphi'(\widehat{\theta}_i; V) = \varphi'(\widehat{\theta}_i), \quad \text{for } i = 0, \dots, k - 1.$$

Note that the ordering of θ_i implies that

$$\prod_{i=0}^{k-1} |\theta_* - \widehat{\theta}_i| \leq \prod_{i=0}^{k-1} |\theta_* - \theta_i|. \quad (17)$$

Since the real analytic function $\varphi'(\theta; V) - \varphi'(\theta)$ is zero at $2k + 1$ distinct points θ_i and $\widehat{\theta}_i$ in $[\theta_* - r, \theta_* + r]$, standard interpolation error results (see, e.g., [6]) yield

$$|\varphi'(\theta; V) - \varphi'(\theta)| = \left| \frac{\varphi^{(2k+2)}(\xi) - \varphi^{(2k+2)}(\xi; V)}{(2k+1)!} \prod_{i=0}^k (\theta - \theta_i) \prod_{i=0}^{k-1} (\theta - \widehat{\theta}_i) \right|, \quad (18)$$

for all $\theta \in [\theta_* - r, \theta_* + r]$ with some $\xi \equiv \xi(\theta) \in [\theta_* - r, \theta_* + r]$. Next, to bound the derivative term, we use that $\varphi(\theta; V)$ extends analytically to the complex disc $D_{\mathbb{C}}(\theta_*; R)$ by Lemma 3. Using Cauchy's integral formula (see, e.g., [21]),

$$|\varphi^{(m)}(\xi; V)| = |\widehat{\varphi}^{(m)}(\xi; V)| = \left| \frac{m!}{2\pi i} \oint_{|z-\theta_*|=R} \frac{\widehat{\varphi}(z; V)}{(z-\xi)^{m+1}} dz \right| \leq \frac{2M \cdot m!}{(R/2)^m}, \quad (19)$$

where the last inequality uses $|\widehat{\varphi}(z; V)| = |\lambda_*(V^* H(z) V)| \leq \|H(z)\|_2 \leq M$, as well as $|z - \xi| \geq |z - \theta_*| - |\theta_* - \xi| \geq R - r \geq R/2$. Setting $V = I_n$, we also obtain

$$|\varphi^{(m)}(\xi)| = |\varphi^{(m)}(\xi; I_n)| \leq \frac{2M \cdot m!}{(R/2)^m}.$$

Plugging these two bounds with $m = 2k + 2$ and (17) into (18) yields (15).

(iic) Because $\varphi'(\theta; V) - \varphi'(\theta)$ has $2k+1$ zeros, Rolle's theorem implies that $\varphi''(\theta; V) - \varphi''(\theta)$ has $2k$ distinct zeros in $[\theta_* - r, \theta_* + r]$. Therefore, in analogy to (15) one can show that $|\varphi''(\theta_*; V) - \varphi''(\theta_*)| \rightarrow 0$ as $r \rightarrow 0$. In turn,

$$|\varphi''(\theta_*) - \varphi''(\theta; V)| \leq |\varphi''(\theta; V) - \varphi''(\theta_*; V)| + |\varphi''(\theta_*; V) - \varphi''(\theta_*)| \rightarrow 0.$$

Because of $\varphi''(\theta_*) < 0$ this shows that (16) is satisfied when $r > 0$ is chosen sufficiently small. \square

Now we are ready to prove the main result.

Proof of Thm. 1. Without loss of generality, we may assume that all θ_i , $i = 0, \dots, k$ are mutually distinct because otherwise $e_{k+1} = 0$ by Lemma 2 and (7) is trivially satisfied.

Let us now assume that r is sufficiently small such that the statements of Lemma 4 and, in turn, Lemma 3 hold. Then $\varphi(\theta; V_k)$ is real analytic on $[\theta_* - R, \theta_* + R]$. The optimality condition $\varphi'(\theta_{k+1}; V_k) = 0$ in the k th iteration of Algorithm 1 combined with Taylor expansion give

$$0 = \varphi'(\theta_{k+1}; V_k) = \varphi'(\theta_*; V_k) + \varphi''(\xi; V_k)(\theta_{k+1} - \theta_*),$$

for some ξ with $|\xi - \theta_*| \leq r$. Therefore,

$$|\theta_{k+1} - \theta_*| = |\varphi'(\theta_*; V_k)| / |\varphi''(\xi; V_k)| \leq 2|\varphi'(\theta_*; V_k)| / |\varphi''(\theta_*)|, \quad (20)$$

where we used (16) in the last inequality. Using the derivative bound (15) and exploiting $\varphi'(\theta_*) = 0$, we obtain

$$|\theta_{k+1} - \theta_*| \leq |\theta_k - \theta_*| \cdot \frac{2C}{|\varphi''(\theta_*)|R} \cdot (k+1) \cdot (2r/R)^{2k}.$$

This implies the monotonicity property (6) when choosing r sufficiently small.

The error recurrence (7) follows immediately from (20), combined with monotonicity (6) and the derivative bound (15). (For notational simplicity, we replaced $R/2$ by R .) \square

3 Crawford number computation

We will now apply and specialize the developments from the previous section to the computation of the Crawford number, which has the form of an eigenvalue optimization problem characterization (2). We apply the subspace method, Algorithm 1, to maximize

$$\varphi(\theta) = \lambda_{\min}(H(\theta)) \quad \text{with} \quad H(\theta) = S \cos \theta + K \sin \theta, \quad (21)$$

with the Hermitian matrices $S = (A + A^*)/2$ and $K = (A - A^*)/2i$. This function is clearly 2π periodic and continuous. When maximizing $\varphi(\theta)$, we can therefore restrict its domain to any interval $[\theta_0, \theta_0 + 2\pi]$.

3.1 Properties of the objective function

Under the assumptions that $\gamma(A) > 0$ and K invertible, a result from [8, Sec. 4.3] shows that, when restricting θ to an appropriately chosen interval of length 2π , the objective function φ is quasi-concave (unimodal) on the open set $\{\theta : \varphi(\theta) > 0\}$. The following lemma strengthens this result, by not requiring invertibility of K and establishing strong concavity.

Lemma 5 (Strong concavity). *Let $\gamma(A) > 0$. Then there exists $\theta_0 \in \mathbb{R}$ such that*

$$\{\theta : \varphi(\theta) > 0\} \cap [\theta_0, \theta_0 + 2\pi] \quad (22)$$

is an open, nonempty interval (ℓ, u) of length at most π . Moreover, $\varphi(\theta)$ is strongly concave on any closed subinterval of (ℓ, u) .

Proof. Let θ_* be a maximizer of $\varphi(\theta)$. By replacing A with $e^{-i\theta_*}A$, we may assume without loss of generality that $\theta_* = 0$, and in turn $\varphi(0) = \gamma(A) > 0$. We will prove the result for $\theta_0 = -\pi$.

We have

$$\varphi(\theta) = \min_{\|v\|_2=1} v^* H(\theta) v = \min_{\|v\|_2=1} v^* S v \cdot \cos \theta + v^* K v \cdot \sin \theta = \min_{x+iy \in \mathcal{F}(A)} x \cos \theta + y \sin \theta. \quad (23)$$

Writing $z = x + iy = |z|(\cos \text{Arg } z + i \sin \text{Arg } z)$ with the argument $\text{Arg } z \in (-\pi, \pi]$, we obtain

$$\varphi(\theta) = \min_{z \in \mathcal{F}(A)} |z| \cdot \cos(\theta - \text{Arg } z). \quad (24)$$

Note that

$$0 < \varphi(0) = \min_{z \in \mathcal{F}(A)} |z| \cos(\text{Arg } z)$$

implies that $\text{Arg } z \in (-\pi/2, \pi/2)$ and, in turn, $(\text{Arg } z - \pi/2, \text{Arg } z + \pi/2) \subset (-\pi, \pi)$ for every $z \in \mathcal{F}(A)$. This allows us to write

$$\begin{aligned} \{\theta : \varphi(\theta) > 0\} \cap [-\pi, \pi] &= \bigcap_{z \in \mathcal{F}(A)} \{\theta : \cos(\theta - \text{Arg } z) > 0\} \cap [-\pi, \pi] \\ &= \bigcap_{z \in \mathcal{F}(A)} (\text{Arg } z - \pi/2, \text{Arg } z + \pi/2). \end{aligned} \quad (25)$$

The right-hand side shows that this set is an interval of length at most π . The left-hand side, together with $\varphi(0) > 0$ and the continuity of $\varphi(\theta)$, shows that this interval is non-empty and open. This proves the first part of the lemma.

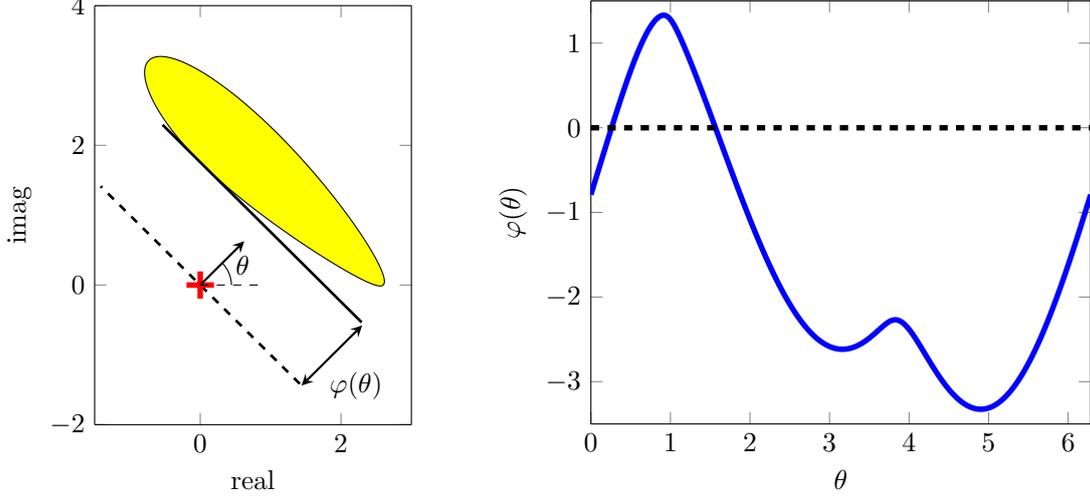


Figure 1: *Left:* Numerical range of a random 4-by-4 matrix (left). *Right:* Corresponding function $\varphi(\theta)$ on the interval $[0, 2\pi]$.

To show strong concavity, let $[\widehat{\ell}, \widehat{u}] \subset (\ell, u)$. Because $\varphi(\theta)$ is continuous there exists $\delta > 0$ such that

$$\varphi(\theta) \geq \delta \quad \text{for } \theta \in [\widehat{\ell}, \widehat{u}].$$

Denoting $\varphi_z(\theta) = |z| \cdot \cos(\theta - \text{Arg } z)$, it follows from equation (24) that

$$\varphi_z''(\theta) = -\varphi_z(\theta) \leq -\min_{z \in \mathcal{F}(A)} \varphi_z(\theta) = -\varphi(\theta) \leq -\delta \quad \text{for } \theta \in [\widehat{\ell}, \widehat{u}].$$

Hence, the function $\varphi(\theta)$ is the minimum of strongly concave functions $\varphi_z(\theta)$ with uniform bound on the second derivative on a closed interval. This proves strong concavity of $\varphi(\theta)$ on $[\widehat{\ell}, \widehat{u}]$. \square

The strong concavity property implies, in particular, that the maximizer θ_* of $\varphi(\theta)$ will be unique within any interval of length 2π when $\gamma(A) > 0$. Moreover, if $\varphi(\theta)$ is twice differentiable at θ_* , we also have

$$\varphi''(\theta_*) < 0. \tag{26}$$

This will for example be the case when $\varphi(\theta_*) = \lambda_{\min}(H(\theta_*))$ is a simple eigenvalue.

Remark 2. *Being continuous and 2π -periodic, the function $\varphi(\theta)$ cannot be strongly concave on an interval of length 2π . On the other hand, Figure 1 (see also [4, Fig. 2]) demonstrates that $\varphi(\theta)$ may well be concave on an interval larger than guaranteed by Lemma 5. In addition, one observes that the negative part of $\varphi(\theta)$ is not necessarily convex.*

As illustrated in Figure 1, the function value $\varphi(\theta)$ admits a geometric interpretation in terms of the numerical range $\mathcal{F}(A)$. Recalling (23), we have

$$\varphi(\theta) = \min_{x+iy \in \mathcal{F}(A)} x \cos \theta + y \sin \theta.$$

This corresponds to the *signed distance* of the numerical range to zero along the direction $\cos \theta + i \sin \theta$ (since the right hand side is the orthogonal projection of points in the numerical range $\mathcal{F}(A)$ onto the direction vector). It also implies that

$$\{x + iy : x \cos \theta + y \sin \theta - \varphi(\theta) = 0\} \tag{27}$$

defines the *supporting line* for the convex set $\mathcal{F}(A) = \mathcal{F}(S + \iota K)$ with normal $\cos \theta + \iota \sin \theta$. Its supporting point is given by

$$p(\theta) = v(\theta)^*(S + \iota K)v(\theta), \quad (28)$$

since the minimum of (23) is attained at an eigenvector $v(\theta)$, as defined in (11). We will use these geometric properties in the following section to get a better understanding of the subspace method.

3.2 Subspace method applied to Crawford number computation

As mentioned before, the full subspace method for computing the Crawford number simply consists of applying Algorithm 1 to the function φ from (21). The objective function of the reduced maximization problem takes the form

$$\varphi(\theta; V_k) = \lambda_{\min}(V_k^* H(\theta) V_k) = \lambda_{\min}(S_k \cos \theta + K_k \sin \theta),$$

where $S_k = V_k^* S V_k = (A_k + A_k^*)/2$ and $K_k = V_k^* K V_k = (A_k - A_k^*)/2\iota$, with $A_k = V_k^* A V_k$. This shows that the reduced problem amounts to computing the Crawford number of A_k :

$$\gamma(A_k) = \max \left\{ \max_{\theta \in [0, 2\pi]} \varphi(\theta; V_k), 0 \right\} \quad (29)$$

Because $\gamma(A_k) \geq \gamma(A)$, one can immediately stop the algorithm and return $\gamma(A) = 0$ once it detects $\gamma(A_k) = 0$.

Since A_k is expected to be a small matrix, the computation of $\gamma(A_k)$, together with the optimal value θ_{k+1} in line 4 of Algorithm 1, is not expensive and can be performed by many algorithms; see, e.g., [11, 8, 26, 3]. Our implementation uses a straightforward adaptation of the criss-cross type algorithm from [20], which we observed to work very well in our numerical experiments.

3.2.1 Properties of the reduction objective function and convergence

Because of (29), the strong concavity property established in Lemma 5 also holds for the reduced objective function $\varphi(\theta; V_k)$. In particular, the maximizer of $\varphi(\theta; V_k)$ is unique (within one period) provided that $\gamma(A) > 0$. This property allows us to gain additional insights into the behavior of the subspace method.

Lemma 6. *Let V be an orthonormal basis of a subspace \mathcal{V} of \mathbb{C}^n . Assume that $\gamma(A) > 0$ and let θ_ℓ, θ_u be such that $\theta_\ell < \theta_u < \theta_\ell + 2\pi$. If $v(\theta_\ell), v(\theta_u) \in \mathcal{V}$, then*

$$\arg \max_{\theta \in [\theta_\ell, \theta_\ell + 2\pi]} \varphi(\theta; V) \in (\theta_\ell, \theta_u) \quad \text{if and only if} \quad \arg \max_{\theta \in [\theta_\ell, \theta_\ell + 2\pi]} \varphi(\theta) \in (\theta_\ell, \theta_u).$$

Proof. Suppose that

$$\theta_V = \arg \max_{\theta \in [\theta_\ell, \theta_\ell + 2\pi]} \varphi(\theta; V) \in (\theta_\ell, \theta_u) \quad \text{and} \quad \theta_* = \arg \max_{\theta \in [\theta_\ell, \theta_\ell + 2\pi]} \varphi(\theta) \notin (\theta_\ell, \theta_u),$$

which implies $\theta_\ell < \theta_V < \theta_u < \theta_* \leq \theta_\ell + 2\pi$. Letting $\beta = \max\{\varphi(\theta_\ell), \varphi(\theta_u), 0\}$, Lemma 1 and Lemma 5 give

$$\varphi(\theta_\ell) = \varphi(\theta_\ell; V) \leq \beta, \quad \varphi(\theta_V; V) > \beta, \quad \varphi(\theta_u) = \varphi(\theta_u; V) \leq \beta, \quad \varphi(\theta_*; V) \geq \varphi(\theta_*) > \beta.$$

In particular, this implies that $\varphi(\theta; V)$ has two different local, strictly positive maxima within an interval of length 2π . This contradicts the strong concavity property from Lemma 5. Hence, $\theta_V \in (\theta_\ell, \theta_u)$ implies $\theta_* \in (\theta_\ell, \theta_u)$. The proof of the other direction proceeds analogously. \square

The particular structure of φ allow us, for the special case of the Crawford number, to remove some of the assumptions of Theorem 1 on the local convergence of the subspace method.

Theorem 3 (Convergence of subspace method). *Consider $A \in \mathbb{C}^{n \times n}$ with $\gamma(A) > 0$. Let θ_* be the maximizer of (21) and assume that $\lambda_{\min}(H(\theta_*))$ is a simple eigenvalue. If the initial guess θ_0 is sufficiently close to θ_* then the iterates $\theta_1, \theta_2, \theta_3, \dots$ produced by Algorithm 1 satisfy the monotonicity property (6) and the error recurrence (7).*

Proof. Lemma 5 shows that the assumption $\varphi''(\theta_*) < 0$ of Theorem 1 is satisfied. It remains to show that the other assumption—all iterates remain close to θ_* —is also satisfied.

Without loss of generality, we may assume that $\theta_0 < \theta_* = 0$ and $\varphi(\theta_0) > 0$. Considering eigenvectors $v(\theta_0), v(\theta_*)$ with $\|v(\theta_0)\|_2 = \|v(\theta_*)\|_2 = 1$, this implies

$$\gamma(A) = v(\theta_*)^* H(\theta_*) v(\theta_*) = v(\theta_*)^* S v(\theta_*), \quad 0 = v(\theta_*)^* H'(\theta_*) v(\theta_*) = v(\theta_*)^* K v(\theta_*). \quad (30)$$

The reduced function during the first iteration takes the form

$$\varphi(\theta; v(\theta_0)) = \lambda_{\min}(v(\theta_0)^* H(\theta) v(\theta_0)) = \alpha \cos \theta + \beta \sin \theta, \quad (31)$$

with $\alpha = v(\theta_0)^* S v(\theta_0)$ and $\beta = v(\theta_0)^* K v(\theta_0)$. Since $\varphi(\theta_0) = \varphi(\theta_0; v(\theta_0)) > 0$ due to Lemma 1, we obtain the super-level set

$$\Omega_0 = \{\theta: \varphi(\theta; v(\theta_0)) \geq \varphi(\theta_0)\} \cap [\theta_0, \theta_0 + 2\pi] = [\theta_0, -\theta_0 + 2 \arctan(\beta/\alpha)]. \quad (32)$$

Because of the monotonicity property in Lemma 1, the iterates of Algorithm 1 satisfy for $k \geq 2$ the inequalities

$$\varphi(\theta_k; v(\theta_0)) \geq \varphi(\theta_k, V_{k-1}) \geq \varphi(\theta_*) \geq \varphi(\theta_0),$$

and, hence, they stay in the superlevel set:

$$\{\theta_0, \theta_1, \dots\} \subset \Omega_0.$$

We now show that the length of Ω_0 becomes arbitrarily small when choosing θ_0 sufficiently close to 0. First, from the interpolation conditions in Lemma 1 it follows

$$\begin{aligned} \varphi(\theta_0) &= \varphi(\theta_0; v(\theta_0)) = \alpha \cos \theta_0 + \beta \sin \theta_0, \\ \varphi'(\theta_0) &= \varphi'(\theta_0; v(\theta_0)) = -\alpha \sin \theta_0 + \beta \cos \theta_0. \end{aligned}$$

By solving this linear equation in $\alpha \equiv \alpha(\theta_0), \beta \equiv \beta(\theta_0)$, we obtain

$$\alpha(\theta_0) = \varphi(\theta_0) \cos \theta_0 - \varphi'(\theta_0) \sin \theta_0, \quad \beta(\theta_0) = \varphi(\theta_0) \sin \theta_0 + \varphi'(\theta_0) \cos \theta_0.$$

Using Taylor expansion around 0 and $\varphi'(0) = 0$, we have

$$\frac{\beta}{\alpha} = \theta_0 + \frac{\varphi''(0)}{\varphi(0)} \theta_0^2 + O(\theta_0^3) \quad (33)$$

and therefore

$$|\Omega_0| = 2 \left| \arctan \frac{\beta}{\alpha} - \theta_0 \right| = 2 \left| \frac{\varphi''(0)}{\varphi(0)} \theta_0 \right|.$$

Since $\varphi(0) = \gamma(A) > 0$ and $\varphi''(0)$ are constants independent of θ_0 , $|\Omega_0| \rightarrow 0$ as $\theta_0 \rightarrow 0$.

In summary, for every $r > 0$, we can attain

$$\{\theta_0, \theta_1, \dots\} \subset \Omega_0 \subset [-r, r],$$

by choosing $\theta_0 < 0$ sufficiently close to 0, which completes the proof. \square

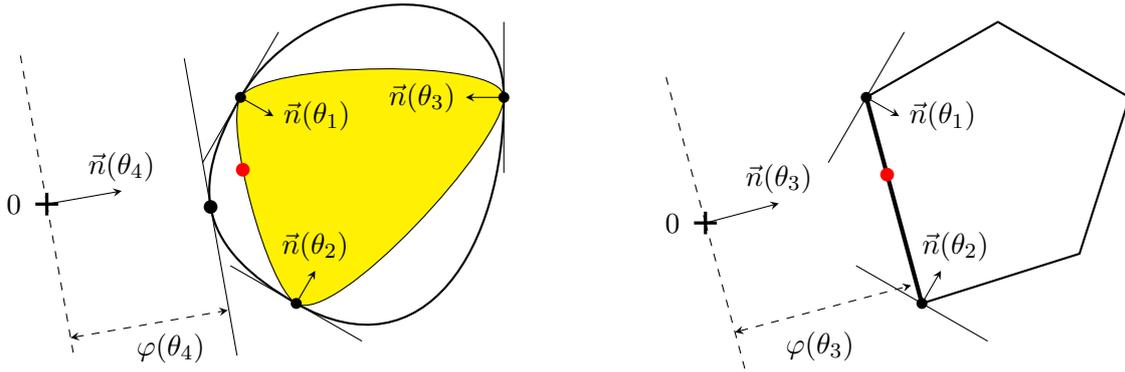


Figure 2: Geometric illustration of the subspace method, Algorithm 1, for Crawford number computation. *Left:* The thick black line and the yellow region correspond to the full and reduced numerical range, respectively. The black dots are the supporting points, and the red dot is the closest point in the yellow region to 0. *Right:* Example of a polygonal numerical range for which the reduced numerical range reduces to a line on the boundary (the thick black line with the red dot).

We end this section with a geometric interpretation of the subspace method in terms of the numerical range; see Figure 2. Recalling (1), we have the reduced Crawford number

$$\gamma(A_k) = \min\{|z|: z \in \mathcal{F}(A_k)\}, \quad A_k = S_k + \iota K_k. \quad (34)$$

It follows immediately from definition (1) that the numerical range of the reduced matrix A_k is a subset of the original one:

$$\mathcal{F}(A_k) = \mathcal{F}(V_k^* A V_k) \subset \mathcal{F}(A).$$

The boundaries of $\mathcal{F}(A)$ and $\mathcal{F}(A_k)$ touch each other at the *supporting points*

$$p(\theta_i) = v_i^*(S + \iota K)v_i, \quad i = 1, \dots, k,$$

because the eigenvectors v_1, \dots, v_k in Algorithm 1 are contained in $\text{span}(V_k)$. The two boundaries also share the supporting lines

$$\{x + \iota y: x \cos \theta_i + y \sin \theta_i - \varphi(\theta_i) = 0\}, \quad i = 1, \dots, k,$$

due to the interpolation $\varphi(\theta_i; V_k) = \varphi(\theta_i)$ together with (27)–(28). In the particular case of a polygonal numerical range (which is equivalent to A being a normal matrix), the reduced numerical range interpolates the edge of the polygon exactly, if the two vertices belong to the supporting points, as illustrated in the right plot of Figure 2.

3.3 The 3-vector subspace method

The concavity property of $\varphi(\theta)$ leading to the bracketing property in Lemma 6 allows us to develop a variant of the subspace method that only uses 3 vectors to generate the projection subspace but still enjoys favorable convergence properties.

Given three sampling points $0 < \ell_k < \theta_k < u_k \leq 2\pi$ such that the global maximum θ_* is contained in $[\ell_k, u_k]$, the k th iteration of our 3-vector subspace method proceeds by first determining an orthonormal basis V_k of the corresponding three eigenvectors $v(\ell_k)$, $v(\theta_k)$, $v(u_k)$. It then

computes

$$\widehat{\theta} = \arg \max_{\theta \in [0, 2\pi]} \lambda_{\min}(V_k^* H(\theta) V_k)$$

By Lemma 6, $\widehat{\theta} \in [\ell_k, u_k]$ and, moreover,

$$\theta_* \in \begin{cases} [\ell_k, \theta_k] & \text{if } \widehat{\theta} \in [\ell_k, \theta_k], \\ [\theta_k, u_k] & \text{otherwise,} \end{cases} \quad (35)$$

$$(36)$$

Therefore, we set $[\ell_{k+1}, \theta_{k+1}, u_{k+1}] = [\ell_k, \widehat{\theta}, \theta_k]$ if $\widehat{\theta} \in [\ell_k, \theta_k]$ and $[\ell_{k+1}, \theta_{k+1}, u_{k+1}] = [\theta_k, \widehat{\theta}, u_k]$ otherwise. In turn, we obtain a shorter interval $[\ell_{k+1}, u_{k+1}]$ that still contains θ_* . This allows us to repeat the process until convergence is achieved, leading to Algorithm 2.

Algorithm 2 3-vector subspace method

Input: Matrix $A \in \mathbb{C}^{n \times n}$, initial guess $\theta_0 \in (0, 2\pi)$, tolerance $\text{tol} > 0$.

Output: Approximation of Crawford number $\gamma(A)$.

- 1: Initialize: $\ell_0 = 0$, $u_0 = 2\pi$, and $V_0 = \text{orth}([v(\ell_0), v(\theta_0), v(u_0)])$.
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: $\theta_{k+1} = \arg \max_{\theta \in [0, 2\pi]} \lambda_{\min}(V_k^* H(\theta) V_k)$ with $\varphi_{k+1} = \lambda_{\min}(V_k^* H(\theta_{k+1}) V_k)$.
 - 4: **if** $\varphi_{k+1} \leq 0$, **then** return $\gamma(A) = 0$.
 - 5: *Stopping criteria:* **if** $k > 0$ **and** $\lambda_k - \varphi_{k+1} < \text{tol} \cdot \varphi_{k+1}$, **then** return $\gamma = \varphi_{k+1}$.
 - 6: *Interval update:*
 if $\theta_{k+1} \in [\ell_k, \theta_k]$, **then** $(\ell_{k+1}, u_{k+1}) = (\ell_k, \theta_k)$, **otherwise** $(\ell_{k+1}, u_{k+1}) = (\theta_k, u_k)$.
 - 7: Compute smallest eigenvalue λ_{k+1} and corresponding eigenvector $v(\theta_{k+1})$ of $H(\theta_{k+1})$.
 - 8: *Subspace update:* $V_{k+1} = \text{orth}([v(\ell_{k+1}), v(\theta_{k+1}), v(u_{k+1})])$
 - 9: **end for**
-

Note that in contrast to the (full) subspace method, Algorithm 2 will not produce the exact solution after $n - 1$ steps. The following theorem shows that it still converges globally and attains fast local convergence.

Theorem 4 (Convergence of 3-vector method). *Let $\gamma(A) > 0$. Then the iterates $\theta_1, \theta_2, \theta_3, \dots$ produced by Algorithm 2 are globally convergent to θ_* . The convergence is locally at least quadratic, provided that $\lambda_{\min}(H(\theta_*))$ is a simple eigenvalue. Moreover, if the sequence is also locally alternating around θ_* , i.e.,*

$$(\theta_{k+1} - \theta_*)(\theta_k - \theta_*) < 0, \quad \text{for all } k > p,$$

with some $p \in \mathbb{N}$, then the local convergence order is improved to $\sigma \approx 2.26953$.

Proof. Since $(\ell_k)_{k=0}^{\infty}$ is monotonically increasing, it converges to a limiting point ℓ . Similarly, $(u_k)_{k=0}^{\infty}$ monotonically decreases to a limiting point u . As (ℓ_k) and (u_k) contain all iterates θ_k for $k = 0, 1, 2, \dots$, it follows that $(\theta_k)_{k=0}^{\infty}$ has at most two limiting points ℓ and u . Let (θ_{k_i}) , $i = 0, 1, \dots$, be a converging subsequence and assume, without loss of generality, that it converges to ℓ .

Since $\text{span}(V_{k_i-1})$ contains $v(\ell_{k_i-1})$, the interpolation and monotonicity properties from Lemma 1 imply

$$\varphi(\ell_{k_i-1}; V_{k_i-1}) = \varphi(\ell_{k_i-1}) \leq \varphi(\theta_*) \leq \max_{\theta \in [0, 2\pi]} \varphi(\theta; V_{k_i-1}) = \varphi(\theta_{k_i}; V_{k_i-1}) \quad (37)$$

and

$$\begin{aligned} |\varphi(\ell_{k_i-1}; V_{k_i-1}) - \varphi(\theta_{k_i}; V_{k_i-1})| &= |\lambda_{\min}(V_{k_i-1}^* H(\ell_{k_i-1}) V_{k_i-1}) - \lambda_{\min}(V_{k_i-1}^* H(\theta_{k_i}) V_{k_i-1})| \\ &\leq \|V_{k_i-1}^* (H(\ell_{k_i-1}) - H(\theta_{k_i})) V_{k_i-1}\|_2. \end{aligned}$$

The right hand side converges to zero due to the continuity of $H(\theta)$, since both subsequences $(\theta_{k_i})_{i=0}^\infty$ and $(\ell_{k_i-1})_{i=1}^\infty$ converge to the same limiting point ℓ . Therefore, the inequalities in (37) become equalities as $k_i \rightarrow \infty$, and it holds that

$$\lim_{k_i \rightarrow \infty} \varphi(\ell_{k_i-1}; V_{k_i-1}) = \lim_{k_i \rightarrow \infty} \varphi(\ell_{k_i-1}) = \varphi(\ell) = \varphi(\theta_*),$$

using the continuity of $\varphi(\theta)$. This shows that any converging subsequence of $(\theta_k)_{k=0}^\infty$ converges to θ_* , and hence $\theta_k \rightarrow \theta_*$.

It remains to discuss the local convergence order. According to the updating strategy in lines 6 and 8, $\text{span}(V_{k+1})$ contains the *two* eigenvectors $v(\theta_{k+1})$ and $v(\theta_k)$. Therefore, θ_{k+2} can be viewed as being produced by the subspace method, Algorithm 1, with a subspace containing $v(\theta_k)$ and $v(\theta_{k+1})$. Note that we may assume $\theta_k \neq \theta_{k+1}$ because otherwise (37) implies $\theta_k = \theta_{k+1} = \theta_*$ and Algorithm 2 returns the exact solution. Due to the local monotonicity behavior of the subspace method, see (6), it holds that $|\theta_{k+2} - \theta_*| \leq |\theta_{k+1} - \theta_*|$ provided that θ_k and θ_{k+1} are sufficiently close to θ_* . Moreover, using the local error recurrence (7), we obtain

$$e_{k+2} \leq 2C \cdot e_{k+1} e_k^2.$$

Then, following the proof of Theorem 2, we obtain a convergence order of $\sigma = 2$, which is the root of $\sigma - \sigma^2 - 2 = 0$ that has magnitude larger than 1.

If the sequence is locally alternating, then the subspace updating scheme in line 6 yields $\{\ell_{k+1}, m_{k+1}, u_{k+1}\} = \{\theta_{k+1}, \theta_k, \theta_{k-1}\}$ for $k > p$. Therefore, the subspace contains the *three* vectors $v(\theta_{k-1})$, $v(\theta_k)$ and $v(\theta_{k+1})$. For the reason mentioned above, we may assume that θ_{k-1} , θ_k , θ_{k+1} are mutually distinct. In turn, the error recurrence (6) improves to

$$e_{k+2} \leq C \cdot e_{k+1} e_k^2 e_{k-1}^2.$$

This leads to a convergence order of about $\sigma = 2.26953$, determined by the root of the cubic equation $\sigma^3 - \sigma^2 - 2\sigma - 2 = 0$ with magnitude larger than 1. \square

Although the 3-vector subspace method enjoys global convergence, the function values φ_k produced by Algorithm 2 are not guaranteed to monotonically decrease because the subspace monotonicity $\text{span}(V_{k+1}) \subset \text{span}(V_k)$ is lost.

Numerically, we usually observed that the alternating property, required for faster local convergence in Theorem 4, holds once θ_k is sufficiently close to θ_* . In the following, we provide some theoretical evidence, by showing that this property holds in the first step. For this purpose, we first establish the following technical result.

Lemma 7. *Let $\gamma(A) > 0$ and let θ_* be the maximizer of (21). Denote $v(\theta)$ as in (11). Assume that $\lambda_{\min}(H(\theta_*))$ is a simple eigenvalue. Then the following statements hold:*

- (a) $\varphi''(\theta_*) \leq -\gamma(A)$;
- (b) $\varphi''(\theta_*) = -\gamma(A)$ if and only if $v(\theta_*)$ is an eigenvector of $H'(\theta_*)$;
- (c) if $v(\theta_*)$ is an eigenvector of $H'(\theta_*)$, then there is $r > 0$ such that $v(\theta) = v(\theta_*)$ for all $|\theta - \theta_*| \leq r$.

Proof. Consider the spectral decomposition $H(\theta_*) = \sum_{i=1}^n \lambda_i v_i v_i^*$ with $\lambda_1 = \lambda_{\min}(H(\theta_*))$ and $v_1 = v(\theta_*)$. Using existing results on the second derivative of eigenvalues [16] and $v_1^* H''(\theta_*) v_1 = -v_1^* H(\theta_*) v_1$, we obtain

$$\varphi''(\theta_*) = v_1^* H''(\theta_*) v_1 + 2 \sum_{i=2}^n \frac{|v_i^* H'(\theta_*) v_1|^2}{\lambda_1 - \lambda_i} = -\gamma(A) - 2 \sum_{i=2}^n \frac{|v_i^* H'(\theta_*) v_1|^2}{\lambda_i - \lambda_1}. \quad (38)$$

This proves (a).

Concerning part (b), (38) shows that $\varphi''(\theta_*) = -\gamma(A)$ if and only if $|v_i^* H'(\theta_*) v_1|^2 = 0$ for $i = 2, \dots, n$. Since $\{v_1, \dots, v_n\}$ is an orthonormal basis, the later condition is equivalent to $H'(\theta_*) v_1 \in \text{span}\{v_1\}$, that is, v_1 is an eigenvector of $H'(\theta_*)$.

To prove part (c), we assume $\theta_* = 0$ without loss of generality (by replacing A with $e^{-i\theta_*} A$). Because of

$$H(\theta) = S \cos \theta + K \sin \theta = \sin(\theta) H'(0) + \cos(\theta) H(0),$$

and $H'(0)v(0) = \alpha v(0)$ for some (fixed) $\alpha \in \mathbb{R}$, we have $H(\theta)v(0) = \tilde{\lambda}v(0)$ with $\tilde{\lambda} = \sin(\theta)\alpha + \cos(\theta)\lambda_1$. Hence $v(0)$ is an eigenvector of $H(\theta)$ with eigenvalue $\tilde{\lambda}$. For sufficiently small $|\theta|$, this remains the smallest eigenvalue of $H(\theta)$ because of the simplicity assumption. \square

The alternating property for the first iterate now follows.

Lemma 8. *Let $\gamma(A) > 0$ and assume that $\lambda_{\min}(H(\theta_*))$ is a simple eigenvalue. Then the first iterate of Algorithm 2 satisfies*

$$(\theta_1 - \theta_*)(\theta_0 - \theta_*) \leq 0, \quad (39)$$

for θ_0 sufficiently close to θ_* .

Proof. Without loss of generality, assume that $\theta_0 < \theta_* = 0$. The first iterate θ_1 is determined as the maximizer of $\varphi(\theta; v(\theta_0))$.

If $v(\theta_*)$ is an eigenvector of $H'(\theta_*)$ then Lemma 7(c) implies $v(\theta_0) = v(\theta_*)$ (provided that θ_0 is sufficiently close to θ_*). Thus, θ_1 is already the exact maximizer and (39) is trivially satisfied. Now, suppose that $v(\theta_*)$ is not an eigenvector of $H'(\theta_*)$ and, hence, $\gamma(A) + \varphi''(\theta_*) < 0$ holds by Lemma 7(a) and (b). Recalling (31) and (33), we obtain from $0 = \varphi'(\theta_1; v(\theta_0))$ that

$$\tan \theta_1 = \frac{\beta}{\alpha} = \frac{\gamma(A) + \varphi''(0)}{\gamma(A)} \theta_0 + O(\theta_0^2).$$

This implies $\theta_1 > 0$ for $\theta_0 < 0$ sufficiently close to 0 and shows (39). \square

4 Numerical experiments

All algorithms discussed in this paper have been implemented in MATLAB¹. The primary purpose of the first experiment is to confirm our theoretical results and to demonstrate that the derived local convergence orders seem to be tight. We also compare with another algorithm [26] for Crawford number computation. In the end, we apply our subspace algorithms to the discretized boundary integral operators from [2], and show the potential of the algorithms for large-scale problems.

¹Available at www.unige.ch/~dlu.

Example 1 Our convergence analysis implies for k sufficiently large that

$$|\varphi_k - \varphi_*| \approx c|\varphi_{k-1} - \varphi_*|^\sigma$$

with $\sigma \approx 2.4142$ for the (full) subspace method (Algorithm 1) and $\sigma \approx 2.2695$ for the 3-vector variant (Algorithm 2). In turn,

$$\ln e_k \approx \sigma \ln e_{k-1} + \ln c \quad \text{with} \quad e_k = |\varphi_k - \varphi_*|.$$

In the following, we numerically verify the value of the slope σ . Since it is rather difficult to test high convergence orders in double precision, we perform the computation with 620 decimal digits using the Advanpix MP toolbox.² We substitute the “exact” value for φ_* by the result of a computation with 700 decimal digits.

Scenario 1: $\lambda_{\min}(H(\theta_*))$ is a simple eigenvalue. Figure 3 plots $\ln e_k$ vs. $\ln e_{k-1}$ for Algorithms 1 and 2 applied to the 120×120 matrix $A = F + \imath M - cI_n$, where F is the Fiedler matrix F , M is the Moler matrix, and $c = 4000 - 4000\imath$, see also [26]. The numerical range of the A is shown in the left plot of Figure 3.

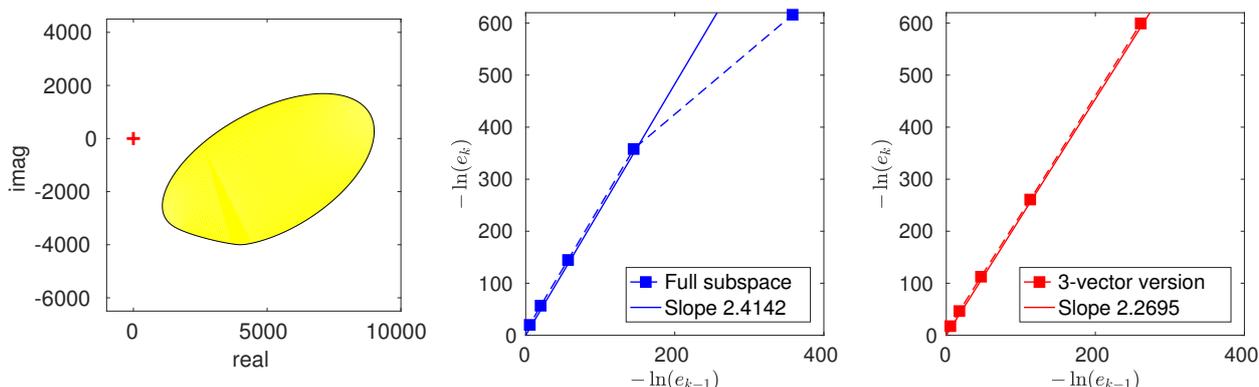


Figure 3: Numerical range (left) and observed order of convergence of Algorithm 1 (middle) and Algorithm 2 (right) for a shifted sum of the Fiedler and Moler matrices.

Figure 4 shows the results obtained for $A = G \cdot \exp(\pi\imath/3) - (4 + 2\imath) \cdot I_n$, where G is the Grcar matrix of size $n = 120$.

For both matrices, we observe that the derived convergence orders from Theorems 3 and 4 seem to be tight.

²Available from www.advanpix.com.

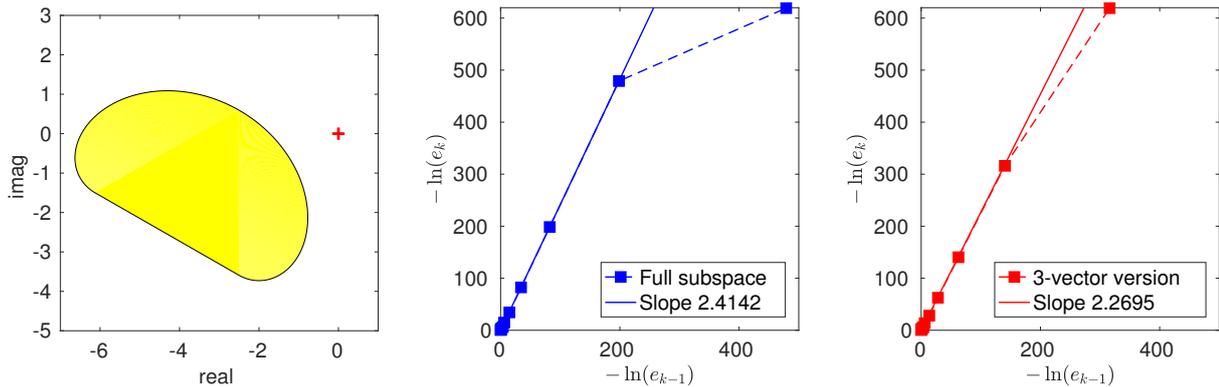


Figure 4: Numerical range (left) and observed order of convergence of Algorithm 1 (middle) and Algorithm 2 (right) for a rotated Grcar matrix.

Scenario 2: $\lambda_{\min}(H(\theta_*))$ is a multiple eigenvalue. Figure 5 shows the convergence of the subspace methods for the tridiagonal matrix

$$A = \text{tridiag} \begin{pmatrix} \imath & \imath & \dots & \imath \\ 1 & 1 & a_3 & \dots & a_n \\ \imath & \imath & \dots & \imath \end{pmatrix} + 0.5\imath \cdot I_n \quad \text{with} \quad a_j = 2 + \frac{j}{n}, \quad (40)$$

with $n = 120$. Visually, the point of the numerical range closest to zero is on the real line. Therefore the eigenvalue relevant for the convergence analysis is the smallest eigenvalue 1 of $(A+A^*)/2$. As this eigenvalue has multiplicity 2, our local convergence results do not apply. Indeed, the convergence plots reveal that the convergence orders established in Theorems 3 and 4 are not attained. On the other hand, one still obtain fast local convergence, seemingly quadratic convergence for the full subspace method and superlinear convergence for the 3-vector variant. Further experiments show that the observed order of convergence remains the same when setting $a_3 = \dots = a_d = 1$, despite the increase of the eigenvalue multiplicity to d . We have also constructed matrices A with $\lambda_{\min}(H(\theta_*))$ a multiple eigenvalue for which the convergence orders were less than the one in Figure 5.

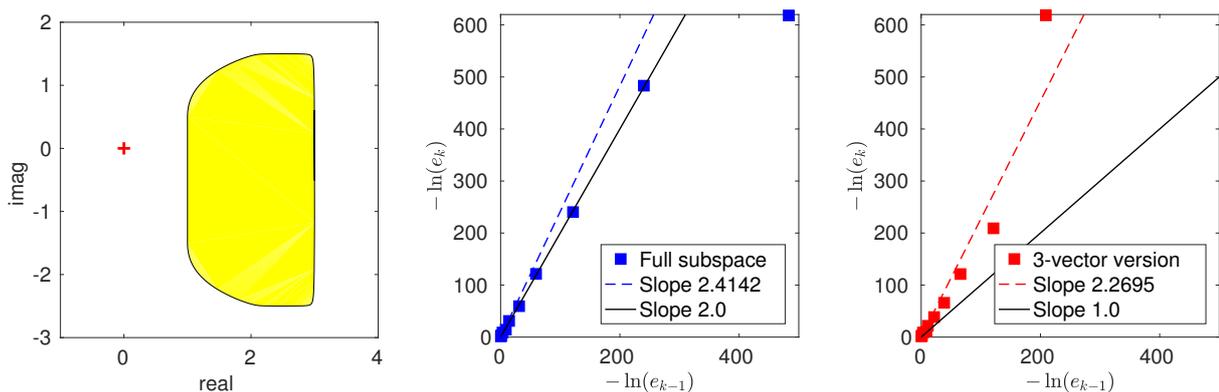


Figure 5: Numerical range (left) and observed order of convergence of Algorithm 1 (middle) and Algorithm 2 (right) for the tridiagonal matrix (40).

Example 2 In this example, we compare the performance of the subspace methods with Uhlig’s algorithm³ [26]. As a test matrix A , we use the Grcar matrix of size $n = 120$ mentioned above, but this time the matrix is not rotated. Although the smallest eigenvalue is not multiple, the small eigenvalues of the symmetric part of A form a cluster. This corresponds to a visually flat vertical portion of the boundary of the numerical range, see Figure 6. In this flat portion, the angle has little impact on the distance which could make it difficult to determine the optimal angle. In our experiment, we compute the Crawford number of the shifted matrices

$$A - (x_0 + y_0i)I_n$$

with $x_0 + y_0i$ marked as black dots in Figure 6. One can equivalently view these shifts as zero and compute its distance to the numerical range.

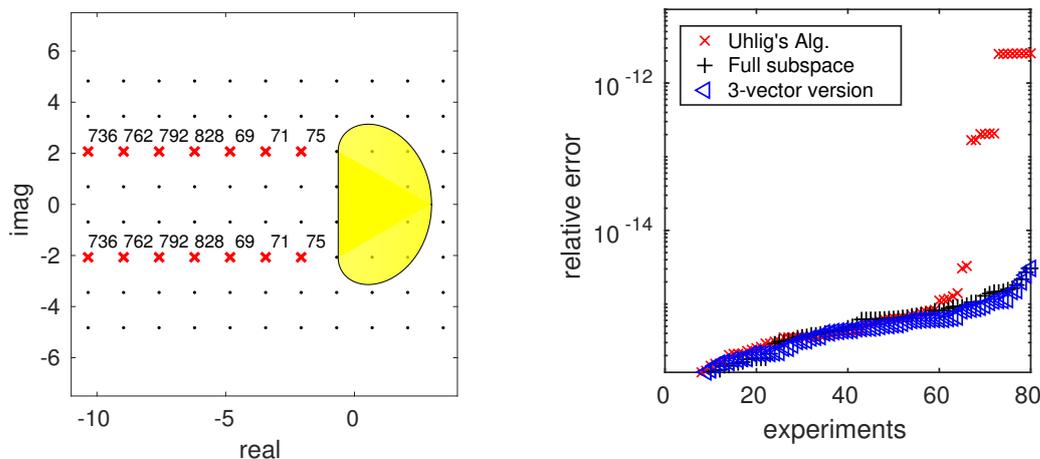


Figure 6: *Left:* Black dots and red crosses denote shifts used in the experiment. Red crosses denote shifts for which Uhlig’s algorithm takes more than 60 iterations to converge, with the number of required iterations marked beside. *Right:* Relative error of the computed Crawford number, sorted in increasing order.

In all experiments, we used the tolerance $\text{tol} = 10^{-13}$ to terminate the algorithms. The iteration numbers shown in Figure 6 show that Uhlig’s algorithm sometimes faces convergence difficulties. In the worst case it requires more than 800 iterations—and so also 800 eigenvalue computations of the matrix $H(\theta)$ —until convergence. In contrast, both Algorithms 1 and 2 converge with a maximum number of 9 iterations and an average of 5.5. The error plot in Figure 6 reveals that the subspace methods are as accurate as Uhlig’s algorithm.⁴

Example 3 We now consider the Crawford number computation problem arising from the study of coercivity constants of boundary integral operators in acoustic scattering [2]. The integral operators of interest are defined by

$$a_k(u, v) = \int_{\Gamma} B_k u(y) \cdot \overline{v(y)} ds(y), \quad \text{with} \quad B_k = I + K_k - ikS_k, \quad (41)$$

³The source code is available from <http://www.auburn.edu/~uhligfd>

⁴For the relative error computation, we compute the ‘accurate’ solution by bisection methods with tolerance equal to the machine precision.

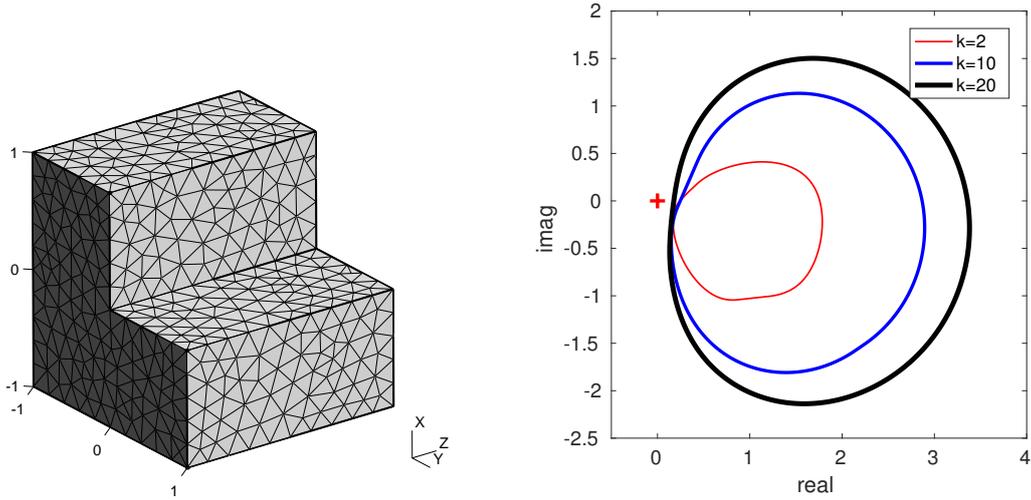


Figure 7: Left: L-shaped obstacle discretized with mesh size $h = 0.2$. Right: Numerical range of the discretized boundary integral operator A_k^h for wave numbers $k = 2, 10, 20$ and $h = 0.2$.

where $k > 0$ is the wave number, Γ is the boundary of a sound-soft bounded obstacle in \mathbb{R}^3 , $u(x), v(x) \in L^2(\Gamma)$, I is the identity operator, and K_k and S_k are defined by

$$K_k u(x) = 2 \int_{\Gamma} \frac{\partial \Phi(x, y)}{\partial n(x)} u(y) ds(y), \quad S_k u(x) = 2 \int_{\Gamma} \Phi(x, y) u(y) ds(y), \quad x \in \Gamma.$$

Here, $\Phi(x, y) = e^{ik|x-y|}/(2\pi|x-y|)$ for $x, y \in \mathbb{R}^3$, $x \neq y$, and $n(x)$ is the outpoint unit normal at Γ . To estimate the coercivity constant γ of the boundary integral operator (41) (i.e., the largest γ with $\gamma \|u\|^2 \leq |a_k(u, u)|$, $\forall u \in L^2(\Gamma)$), we discretize it to a matrix A_k^h and compute its Crawford number. In order to achieve a good estimation, it is preferable to use a small mesh size h in discretization. This in general leads to a large dense matrix A_k^h for which the eigenvalue computation of $H(\theta)$ is expensive. So it is crucial to compute the Crawford number using as few eigenvalue evaluations as possible.

In our experiment, we consider the 3D time-harmonic acoustic scattering from an L-shaped sound-soft bounded obstacle $\Gamma \subset \mathbb{R}^3$ displayed in Figure 7. We discretize the boundary integral operators (41) by the Galerkin boundary element library BEM++⁵, with triangular mesh (generated by gmsh⁶ with Delaunay's algorithm) and piecewise linear basis functions. For a detailed description of the Galerkin discretization of the boundary integral operator, we refer to [23]. In order to speed up the assembly of the boundary integral operator, BEM++ stores the coefficient matrix A_k^h as a hierarchical matrix [9]. Since matrix vector and transpose matrix vector products can be efficiently evaluated for hierarchical matrices, we compute the smallest eigenvalue of $H(\theta)$ using the MATLAB function `eigs` [17].

We summarize in Table 1 the experimental results for three wave numbers $k = 2, 10, 20$. The boundary integral operators were discretized with mesh size $h = 0.02$ to matrices A_k^h of size $n = 396162$. From the timing statistics, it is clear that the subspace acceleration Algorithm 1 has

⁵The software is available from <http://www.bempp.org/>

⁶The software is available from <http://gmsh.info/>

significantly reduced the computational cost. As in the previous experiment, both Algorithm 1 and Uhlig’s algorithm terminate with tolerances $\text{tol} = 10^{-13}$.

Table 1: Computation results for Example 3 with mesh size $h = 0.02$.

wave number k	memory A_k^h		Crawford number $\gamma(A_k^h)$	its.	timing (h)
2	46 GB	Alg. 1	1.556145884392413e-01	5	2.5
		Uhlig’s	1.556145884392416e-01	11	4.4
10	54 GB	Alg. 1	1.880394259192281e-01	7	6.1
		Uhlig’s	1.880394259192323e-01	30	29.9
20	60 GB	Alg. 1	1.777716873410842e-01	8	8.1
		Uhlig’s	1.777716873410810e-01	24	24.9

5 Conclusions

We have analyzed the convergence of the subspace method for univariate eigenvalue optimization. The obtained convergence order not only improves upon existing results but it also appears to be tight. For the special case of the Crawford number, we have established novel properties of the objective function. This has resulted in a three-dimensional subspace method that is proven to enjoy favorable convergence properties and seems to work very well for the examples considered in this paper.

The developments in this paper offer several possibilities for future research. In particular, this concerns the proper treatment of large-scale problems and computing variations of the Crawford number, such as the inner numerical radius [4]. In addition, the subspace acceleration and its bracketing property may also prove useful for the shrinking problem in [10].

References

- [1] N. Aliyev, P. Benner, E. Mengi, and M. Voigt. Large-scale computation of \mathcal{H}_∞ -infinity norms by a greedy subspace method. Technical report, 2016.
- [2] T. Betcke and E. A. Spence. Numerical estimation of coercivity constants for boundary integral operators in acoustic scattering. *SIAM J. Numer. Anal.*, 49(4):1572–1601, 2011.
- [3] S. Bora and R. Srivastava. Distance problems for Hermitian matrix pencils with eigenvalues of definite type. *SIAM J. Matrix Anal. Appl.*, 37(1):53–85, 2016.
- [4] S. H. Cheng and N. J. Higham. The nearest definite pair for the Hermitian generalized eigenvalue problem. *Linear Algebra Appl.*, 302:63–76, 1999.
- [5] C. R. Crawford. A stable generalized eigenvalue problem. *SIAM J Numer. Anal.*, 13(6):854–860, 1976.
- [6] C. De Boor. *A Practical Guide to Splines*, volume 27. Springer-Verlag New York, 1978.

- [7] C.-H. Guo, N. J Higham, and F. Tisseur. Detecting and solving hyperbolic quadratic eigenvalue problems. *SIAM Journal on Matrix Analysis and Applications*, 30(4):1593–1613, 2009.
- [8] C.-H. Guo, N. J. Higham, and F. Tisseur. An improved arc algorithm for detecting definite hermitian pairs. *SIAM J. Matrix Anal. Appl.*, 31(3):1131–1151, 2009.
- [9] W. Hackbusch. *Hierarchical Matrices: Algorithms and Analysis*. Springer, 2015.
- [10] N. J. Higham, N. Strabac, and V. Sego. Restoring definiteness via shrinking, with an application to correlation matrices with a fixed block. *SIAM Review*, 58(2):245–263, 2016.
- [11] N. J. Higham, F. Tisseur, and P. M. Van Dooren. Detecting a definite hermitian pair and a hyperbolic or elliptic quadratic eigenvalue problem, and associated nearness problems. *Linear Algebra Appl.*, 351:455–474, 2002.
- [12] W. Kahan. Spectra of nearly Hermitian matrices. *Proceedings of the American Mathematical Society*, 48(1):11–17, 1975.
- [13] F. Kangal, K. Meerbergen, E. Mengi, and W. Michiels. A subspace method for large scale eigenvalue optimization. *arXiv preprint arXiv:1508.04214*, 2015.
- [14] T. Kato. *Perturbation Theory for Linear Operators*, volume 132. Springer Science & Business Media, 1995.
- [15] D. Kressner and B. Vandereycken. Subspace methods for computing the pseudospectral abscissa and the stability radius. *SIAM J. Matrix Anal. Appl.*, 35(1):292–313, 2014.
- [16] P. Lancaster. On eigenvalues of matrices dependent on a parameter. *Numer. Math.*, 6(1):377–387, 1964.
- [17] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK users’ guide: solution of large-scale eigenvalue problems with Implicitly Restarted Arnoldi Methods*. SIAM, 1997.
- [18] D. Lu and B. Vandereycken. Criss-cross type algorithms for computing the real pseudospectral abscissa. submitted, 2016.
- [19] K. Meerbergen, E. Mengi, W. Michiels, and R. Van Beeumen. Computation of pseudospectral abscissa for large scale nonlinear eigenvalue problems. *IMA J. Numer. Anal.*, 2016. To appear.
- [20] 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.
- [21] W. Rudin. *Real and Complex Analysis*. Tata McGraw-Hill Education, 1987.
- [22] P. Sirković and D. Kressner. Subspace acceleration for large-scale parameter-dependent Hermitian eigenproblems. *SIAM J Matrix Anal. Appl.*, 37(2):695–718, 2016.
- [23] W. Śmigaj, T. Betcke, S. Arridge, J. Phillips, and M. Schweiger. Solving boundary integral problems with BEM++. *ACM Transactions on Mathematical Software (TOMS)*, 41(2):6, 2015.
- [24] G. W. Stewart. Perturbation bounds for the definite generalized eigenvalue problem. *Linear Algebra Appl.*, 23:69–85, 1979.
- [25] G. W. Stewart and J. Sun. *Matrix Perturbation Theory*. Academic Press, 1990.

- [26] F. Uhlig. On computing the generalized Crawford number of a matrix. *Linear Algebra Appl.*, 438(4):1923–1935, 2013.