

# An Experimental Comparison of Methods for Computing the Numerical Radius

Tim Mitchell\*      Michael L. Overton†

October 6, 2023

## Abstract

We make an experimental comparison of methods for computing the numerical radius of an  $n \times n$  complex matrix, based on two well-known characterizations, the first a nonconvex optimization problem in one real variable and the second a convex optimization problem in  $n^2 + 1$  real variables. We make comparisons with respect to both accuracy and computation time using publicly available software.

AMS Subject Classifications: 15A60, 90C22

## 1 Introduction

Let  $A \in \mathbf{M}_n$ , the space of  $n \times n$  complex matrices. The numerical radius  $r(A)$  is defined as the maximum, in modulus, of the points in the field of values (numerical range), i.e.,

$$r(A) = \max \{ |v^* A v| : v \in \mathbb{C}^n, \|v\|_2 = 1 \}.$$

A well-known result (Johnson [Joh78]; see also Kippenhahn [Kip51]) is

$$r(A) = \max_{\theta \in [0, 2\pi)} h(\theta) \tag{1}$$

where

$$h(\theta) = \lambda_{\max}(H(\theta))$$

and

$$H(\theta) = \frac{1}{2} \left( e^{i\theta} A + e^{-i\theta} A^* \right),$$

with  $\lambda_{\max}(\cdot)$  denoting largest eigenvalue of a Hermitian matrix,  $i$  the imaginary unit and  $^*$  denoting complex conjugate transpose. Although the optimization problem (1) has only one real variable, it is nontrivial to solve because  $h$  is nonconvex. Various algorithms have been introduced to solve (1), including the level-set method of Mengi and Overton [MO05] (based on earlier work of He and Watson [HW97] and of Boyd and Balakrishnan [BB90]), and a cutting plane method of Uhlig [Uhl09] (based on earlier work of Johnson [Joh78]). More recently, improvements to both of these approaches have been given by Mitchell [Mit23], along with a state-of-the-art hybrid method that combines the virtues of both in order to remain efficient in all cases.

---

\*Department of Computer Science, Queens College/CUNY, [tim.mitchell@qc.cuny.edu](mailto:tim.mitchell@qc.cuny.edu)

†Courant Institute of Mathematical Sciences, New York University, [mo1@nyu.edu](mailto:mo1@nyu.edu)

A second well-known characterization of  $r(A)$  (Mathias [Mat93]; see also Ando [And73]) is

$$r(A) = \min_{c \in \mathbb{R}, Z \in \mathbf{M}_n, Z = Z^*} \left\{ c : \begin{bmatrix} cI_n + Z & A \\ A^* & cI_n - Z \end{bmatrix} \succeq 0 \right\}, \quad (2)$$

where  $M \succeq 0$  means a Hermitian matrix  $M$  is positive semidefinite. Thus,  $r(A)$  can be computed by solving a convex optimization problem, known as a semidefinite program (SDP), in  $n^2 + 1$  real variables. If  $A$  is real, then  $Z$  can be restricted to be real symmetric, so the optimization problem has  $n(n+1)/2 + 1$  variables. It is well known that such convex optimization problems can be solved, up to any given accuracy, in polynomial time using the Turing complexity model; the traditional method guaranteeing such complexity is the ellipsoid method, but more practical interior-point methods have also recently been designed with polynomial-time complexity [dKV16].

However, for all practical purposes, interior-point methods for computing  $r(A)$  via (2) are very slow compared to more direct methods based on (1). In this note we illustrate this with some experimental computations carried out using Mitchell's codes available from the supplemental materials published in [Mit23] as well as the latest stable release (version 2.2, build 1148) of the well-known CVX software for "disciplined convex programming" [GB14]. CVX uses a variety of primal-dual interior-point "solvers" for SDP; we use two of these, SDPT3 and SeDuMi. CVX lacks the polynomial-time guarantees of the methods described by [dKV16], but it is much faster than such methods, partly because it uses floating point arithmetic, while any rigorous polynomial-time method would necessarily have to be implemented using arbitrary precision integer or rational arithmetic, as discussed at the end of [dKV16].

## 2 The experiments and the results

We experimented with the following methods using MATLAB R2023a Update 4 and a 2023 Mac Mini configured with a 12-core M2 Pro Apple Silicon CPU and 16GB RAM running macOS 13.5:

- LSO: solve (1) via the Level Set with Optimization Algorithm 3.1 in [Mit23]<sup>1</sup>
- CF: solve (1) using Chebfun via adaptive Chebyshev interpolation<sup>2</sup>
- SDPT3: solve (2) using CVX [GB14] with the SDPT3 solver
- SeDuMi: solve (2) using CVX [GB14] with the SeDuMi solver

We conducted our experiments using randomly generated real as well as complex matrices ranging in size from  $n = 10$  to  $n = 800$ ; the matrix coefficients (real parts and imaginary parts, as appropriate) were generated using the standard normal distribution. Figure 1 shows the running times for our test matrices, with the left panel showing the results for real matrices and the right panel the results for complex matrices. We see immediately that LSO is far more efficient than the other three methods, and that among the other three, CF is the fastest, then SDPT3, and then SeDuMi.<sup>3</sup> Meanwhile, Figure 2 shows, for each pairing of the four methods, the relative discrepancies between the computed values of the numerical radius  $r$ . We see that methods LSO and CF have relative discrepancies always near  $10^{-15}$ , indicating that these methods compute  $r$  to about 14 correct digits, quite impressive considering that MATLAB's IEEE double format

<sup>1</sup>We do not include Mitchell's improved cutting-plane method, Algorithm 5.1 in [Mit23], or his hybrid method that combines Algorithms 3.1 and 5.1, as it is not necessary in our benchmark. Although Algorithm 5.1 can be many times faster than the Algorithm 3.1 on certain problems and many times slower on others, on the random matrices that we test here, Algorithm 3.1 and Algorithm 5.1 generally perform similarly.

<sup>2</sup>Chebfun is a package for globally approximating functions to machine precision accuracy using adaptive Chebyshev interpolation techniques; for more details, see [DHT14]. By method CF, we mean using Chebfun to accurately approximate  $h$  on  $[0, 2\pi]$  followed by a call to the Chebfun `max` routine to find the maximal value, as proposed in [Uhl09].

<sup>3</sup>We note that CVX uses precompiled binary files for both SDPT3 and SeDuMi; if these were instead implemented in MATLAB, they'd presumably be significantly slower.

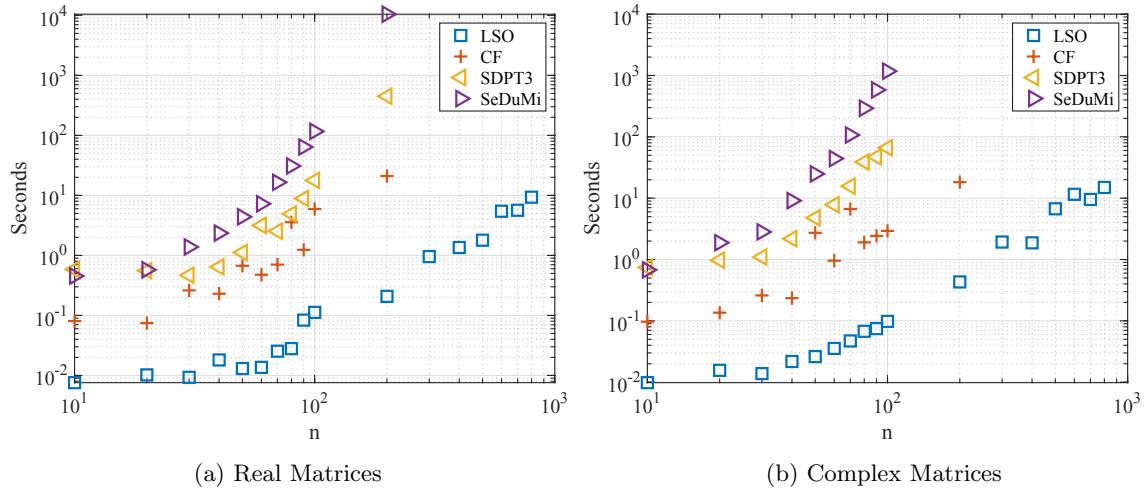


Figure 1: The elapsed running times of the four methods tested are shown. For problems that took at most 3 minutes to solve, the reported times are the average of five trials in order to account for timing variability, while all other reported times are from a single run. Due to their high cost, CF, SDPT3, and SeDuMi were only tested up to size  $n = 200$ , and on the complex matrix with  $n = 200$ , both SDPT3 and SeDuMi crashed, and thus their times are only reported up to  $n = 100$  in the right plot.

numbers have only 53 bits of precision, or approximately 16 decimal digits. The discrepancies between the LSO method and the CVX methods, as well as between the CF method and the CVX methods, indicate that the CVX methods solving (2) return results which are accurate to only about 9 to 12 digits, although we note that CVX was instructed to compute the results with high precision accuracy.<sup>4</sup>

### 3 Conclusion

Our experiments show that although the characterization of the numerical radius as the solution of a semidefinite program is a valuable theoretical tool, as a practical matter, Mitchell's Algorithm 3.1 based on the nonconvex characterization is much faster. We note, however, that the SDP formulation has the appealing property that it also allows the convenient solution of linearly parameterized optimization problems involving the numerical radius, such as the one given in [LO20, Sec. 2.1].

### References

- [And73] T. Ando. Structure of operators with numerical radius one. *Acta Sci. Math. (Szeged)*, 34:11–15, 1973.
- [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  $L_\infty$ -norm. *Systems Control Lett.*, 15(1):1–7, 1990.

<sup>4</sup>By setting `cvx_precision high` in the calling code.

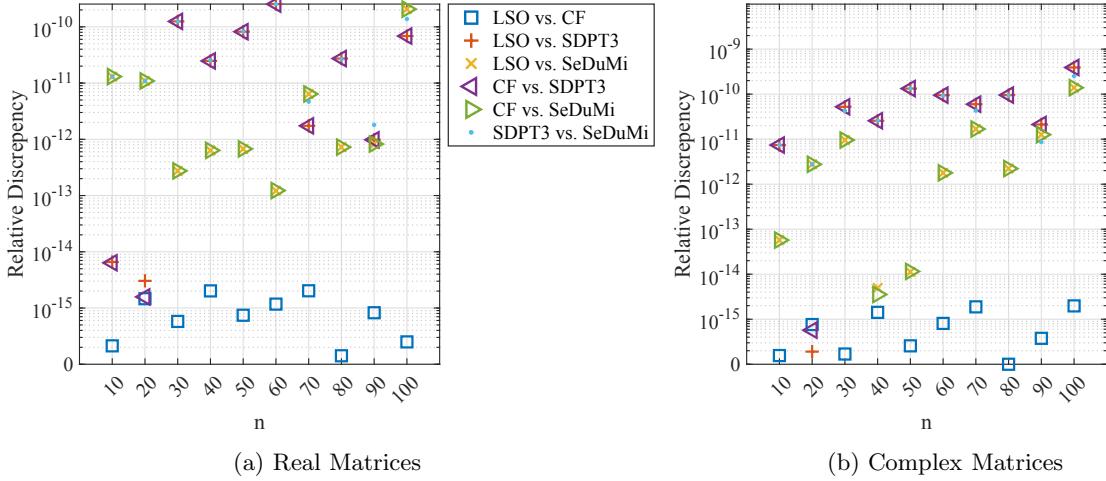


Figure 2: The plots give the pairwise relative discrepancies of the methods' computed numerical radius estimates, i.e.,  $|(r_1(A) - r_2(A))| / \max(r_1(A), r_2(A))$ , where  $r_1(A)$  and  $r_2(A)$  are the computed results for the two methods being compared on matrix  $A$ .

- [DHT14] T. A Driscoll, N. Hale, and L. N. Trefethen. *Chebfun Guide*. Pafnuty Publications, Oxford, UK, 2014.
- [dKV16] Etienne de Klerk and Frank Vallentin. On the Turing model complexity of interior point methods for semidefinite programming. *SIAM J. Optim.*, 26(3):1944–1961, 2016.
- [GB14] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [HW97] C. He and G. A. Watson. An algorithm for computing the numerical radius. *IMA J. Numer. Anal.*, 17(3):329–342, June 1997.
- [Joh78] Charles R. Johnson. Numerical determination of the field of values of a general complex matrix. *SIAM J. Numer. Anal.*, 15(3):595–602, 1978.
- [Kip51] Rudolf Kippenhahn. Über den Wertevorrat einer Matrix. *Math. Nachr.*, 6:193–228, 1951.
- [LO20] A. S. Lewis and M. L. Overton. Partial smoothness of the numerical radius at matrices whose fields of values are disks. *SIAM J. Matrix Anal. Appl.*, 41(3):1004–1032, 2020.
- [Mat93] Roy Mathias. Matrix completions, norms and Hadamard products. *Proc. Amer. Math. Soc.*, 117(4):905–918, 1993.
- [Mit23] Tim Mitchell. Convergence rate analysis and improved iterations for numerical radius computation. *SIAM J. Sci. Comput.*, 45(2):A753–A780, 2023.
- [MO05] Emre Mengi and Michael 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.
- [Uhl09] Frank Uhlig. Geometric computation of the numerical radius of a matrix. *Numer. Algorithms*, 52(3):335–353, 2009.