



Local minimizers of the Crouzeix ratio: a nonsmooth optimization case study

Michael L. Overton¹

Received: 29 May 2021 / Revised: 29 May 2021 / Accepted: 29 October 2021

© The Author(s) under exclusive licence to Istituto di Informatica e Telematica (IIT) 2021

Abstract

Given a square matrix A and a polynomial p , the Crouzeix ratio is the norm of the polynomial on the field of values of A divided by the 2-norm of the matrix $p(A)$. Crouzeix's conjecture states that the globally minimal value of the Crouzeix ratio is 0.5, regardless of the matrix order and polynomial degree, and it is known that 1 is a frequently occurring locally minimal value. Making use of a heavy-tailed distribution to initialize our optimization computations, we demonstrate for the first time that the Crouzeix ratio has many other locally minimal values between 0.5 and 1. Besides showing that the same function values are repeatedly obtained for many different starting points, we also verify that an approximate nonsmooth stationarity condition holds at computed candidate local minimizers. We also find that the same locally minimal values are often obtained both when optimizing over real matrices and polynomials, and over complex matrices and polynomials. We argue that minimization of the Crouzeix ratio makes a very interesting nonsmooth optimization case study, illustrating among other things how effective the BFGS method is for nonsmooth, nonconvex optimization. Our method for verifying approximate nonsmooth stationarity is based on what may be a novel approach to finding approximate subgradients of max functions on an interval. Our extensive computations strongly support Crouzeix's conjecture: in all cases, we find that the smallest locally minimal value is 0.5.

Keywords Field of values · Numerical range · Crouzeix conjecture · Nonsmooth optimization · Chebfun · BFGS

Mathematics Subject Classification 15A60 · 49J52

✉ Michael L. Overton
mol@nyu.edu

¹ Courant Institute of Mathematical Sciences, New York University, New York, United States

1 Introduction

Let \mathcal{P}_m denote the space of polynomials with complex coefficients and degree at most m , let \mathcal{M}^n denote the space of $n \times n$ complex matrices, and, for $A \in \mathcal{M}^n$, let $W(A)$ denote the field of values (numerical range) of A , namely

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

The field of values is a convex, compact set in the complex plane [14]. For $p \in \mathcal{P}_m$ and $A \in \mathcal{M}^n$, the Crouzeix ratio is defined to be

$$f(p, A) = \frac{\|p\|_{W(A)}}{\|p(A)\|_2} \quad (2)$$

where the numerator is

$$\max \{|p(z)| : z \in W(A)\} \quad (3)$$

and the denominator is the 2-norm of the matrix $p(A)$. Consider the optimization problem

$$\min_{p \in \mathcal{P}_m, A \in \mathcal{M}^n} f(p, A).$$

This is a nonconvex, nonsmooth optimization problem: at some pairs (p, A) , f is not differentiable. Crouzeix's famous 2004 conjecture in matrix theory [7] postulates that the globally minimal value of f is 0.5, regardless of n and m . It was established in 2017 [5] that the globally minimal value is no less than $\sqrt{2} - 1 \approx 0.414$.

In our previous work with A. Greenbaum [11], we reported on computational results minimizing f over $n \times n$ real Hessenberg matrices A and real polynomials p with degree at most $n - 1$, for $n = 3$ through $n = 8$. Using the BFGS method, we ran 100 optimization runs for each value of n , initialized with the entries of A and the coefficients of p set randomly using the standard normal distribution. We found that almost all optimization runs generated function values converging either to the conjectured globally minimal value 0.5, or to the evidently locally minimal value 1. In the former case, making use of the generalized null space decomposition [13], we confirmed that the computed final (p, A) always approximated the conjectured globally minimal ‘‘Crabb matrix’’ configurations, to be described below, with field of values a circular disk. In the latter case, using the Schur factorization, we confirmed that the final (p, A) always approximated an ‘‘ice-cream-cone’’ configuration, again to be described below. However, we were not sure whether other locally optimal values of the Crouzeix ratio besides 0.5 and 1 might exist, writing ‘‘...for $n = 7$ and 8, restarting BFGS at and near the final computed pairs [for which f was not close to 0.5 or 1] did not lead to much improvement, suggesting the possibility that there are other stationary values of f between 0.5 and 1’’.

In this paper, we report on much more extensive computations minimizing f , presenting clear evidence that, in fact, there are many other stationary values of f , specifically locally minimal values of f , between 0.5 and 1. Our new computational results also represent perhaps the strongest evidence yet that Crouzeix's conjecture is true. We

also think that they present an interesting “case study” in nonsmooth optimization. In particular, they illustrate how effective the BFGS method is at finding locally minimal values even when these occur at nonsmooth stationary points, despite the fact that the method uses only function and gradient information, not subgradient information. They also illustrate how, once candidate stationary points have been identified, approximate stationarity can be verified numerically even at nonsmooth stationary points, using approximate subgradient information computed at these points.

The paper is organized as follows. In the next section, we explain how nonsmoothness of the Crouzeix ratio f arises. In Sect. 3 we describe the Crabb matrix and ice-cream-cone configurations associated with the stationary values 0.5 and 1 respectively. We briefly describe the computational model in Sect. 4 and explain how to verify approximate nonsmooth stationarity in Sect. 5. We present our experimental results in Sect. 6, and make some concluding remarks in Sect. 7.

2 Nonsmoothness of the Crouzeix ratio

The Clarke subdifferential (generalized gradient) [4] of a locally Lipschitz function h mapping a Euclidean space to \mathbb{R} , evaluated at a point x , is the convex hull of the gradient limits

$$\partial h(x) = \text{conv} \left\{ \lim_{x_k \rightarrow x} \nabla h(x_k) \right\},$$

where the limit is taken over all sequences (x_k) converging to x on which h is differentiable. Elements of $\partial h(x)$ are called (Clarke) subgradients. Clearly, if h is continuously differentiable at x , then $\partial h(x) = \{\nabla h(x)\}$. If $0 \in \partial h(x)$, we say that x is a (Clarke) stationary point of h . If, in addition, h is differentiable at x with $\partial h(x) = \{\nabla h(x)\}$, we say that x is a smooth stationary point; otherwise, it is a nonsmooth stationary point.

By identifying $p \in \mathcal{P}_m$ with its vector of coefficients $[c_0, \dots, c_m]^T \in \mathbb{C}^{m+1}$, we can view the Crouzeix ratio f given in (2) as a function mapping the Euclidean space $\mathbb{E} = \mathbb{C}^{m+1} \times \mathcal{M}^n$, with real inner product

$$\langle (c, A), (d, B) \rangle = \text{Re}(c^* d + \text{tr}(A^* B)),$$

to \mathbb{R} , where $*$ denotes complex conjugate transpose. The function f is locally Lipschitz on the set of pairs (p, A) for which $p(A)$ is not zero.

By the maximum modulus theorem, the maximum value of $|p(z)|$ on $W(A)$ must be attained on $\text{bd } W(A)$, the boundary of the field of values—and only there, unless p is constant on $W(A)$, which can only occur if p is constant or A is a multiple of the identity matrix, cases that are of no interest. The most important source of nonsmoothness of f is that

$$Z(p, A) = \{z \in \text{bd } W(A) : \|p\|_{W(A)} = |p(z)|\} \tag{4}$$

may contain multiple points. In particular, this occurs at the conjectured global minimizers described in the next section.

As explained in [11], there are two other possible sources of nonsmoothness in f . One possibility is that even if $\|p\|_{W(A)}$ is attained only at a single point $z \in \text{bd } W(A)$, the equation $z = v^*Av$ in (1) holds for two or more linearly independent unit vectors v . The other possibility is that the maximum singular value of $p(A)$, which defines the denominator of the Crouzeix ratio, has multiplicity two or more. Since neither of these cases occurs either at the conjectured global minimizers or at the apparent local minimizers found in our computations, we will not consider them further.

3 The Crabb matrix, conjectured global minimizers, and ice-cream-cone stationary points

Pairs (\tilde{p}, \tilde{A}) for which the Crouzeix ratio is 0.5 are known. Given an integer k with $2 \leq k \leq \min(n, m + 1)$, define the polynomial $\tilde{p} \in \mathcal{P}_m$ by $\tilde{p}(\zeta) = \zeta^{k-1}$, set the matrix $\Xi_k \in \mathcal{M}^k$ to

$$\begin{bmatrix} 0 & 2 \\ 0 & 0 \end{bmatrix} \text{ if } k = 2, \text{ or } \begin{bmatrix} 0 & \sqrt{2} & & & & \\ & \cdot & 1 & & & \\ & & \cdot & \cdot & & \\ & & & \cdot & \cdot & \\ & & & & \cdot & 1 \\ & & & & & \cdot & \sqrt{2} \\ & & & & & & 0 \end{bmatrix} \text{ if } k > 2, \tag{5}$$

and set $\tilde{A} = \text{diag}(\Xi_k, 0) \in \mathcal{M}^n$. The matrix Ξ_k was called the Choi-Crouzeix matrix of order k in [11], but after the paper was published, A. Salemi¹ informed us that it was introduced much earlier in a different context by Crabb [6]. The field of values of the Crabb matrix Ξ_k is the unit disk, so the numerator of the Crouzeix ratio is 1, and $\tilde{p}(\tilde{A}) = \tilde{A}^{k-1} = \text{diag}(\Xi_k^{k-1}, 0)$ is a matrix with just one nonzero, namely a 2 in the $(1, k)$ position, so the denominator is 2; hence, the ratio is 0.5.

Since $|\tilde{p}|$ is constant on the boundary of the unit disk, we have that

$$Z(\tilde{p}, \tilde{A}) = \{z \in \mathbb{C} : |z| = 1\},$$

the unit circle, resulting in nonsmoothness of the Crouzeix ratio f at (\tilde{p}, \tilde{A}) . Together with A. Greenbaum and A.S. Lewis [10], we derived the Clarke subdifferential of f at (\tilde{p}, \tilde{A}) , and established that (\tilde{p}, \tilde{A}) is a nonsmooth stationary point of f . The analysis also showed that f is directionally differentiable at (\tilde{p}, \tilde{A}) and that the directional derivative of f is nonnegative in every direction in \mathbb{E} . Although this does not imply that (\tilde{p}, \tilde{A}) is even a local minimizer of f , a discussion of how one might extend the

¹ Private communication, 2017

analysis towards the goal of proving local (but not global!) minimality is given in [10, p. 242].

The property that $f(\tilde{p}, \tilde{A}) = 0.5$ extends easily to pairs (p, A) where $p(\zeta) = (\zeta - \lambda)^{k-1}$ and

$$A = \lambda I + \beta U \text{diag}(\Xi_k, B) U^*,$$

for any $\lambda \in \mathbb{C}$, nonzero $\beta \in \mathbb{C}$, unitary matrix U , and matrix B with $W(B)$ contained in the unit disk. We conjecture that such pairs (p, A) are the *only* ones for which $f(p, A) = 0.5$. However, if the condition that p is a polynomial is relaxed to allow it to be any analytic function, there are many choices for (p, A) for which the ratio is 0.5; for the case $n = 3$, see [8, Sec. 10].

In our computations, we often find the locally minimal value 1. This occurs when the matrix A is block diagonal of the form $A = \text{diag}(\lambda, B)$, with $\lambda \in \mathbb{C}$, $\lambda \notin W(B)$. In this configuration, $W(A) = \text{conv}(\lambda, W(B))$ with $\text{bd } W(A)$ consisting only of λ , part of $\text{bd } W(B)$ and two line segments connecting λ to $W(B)$. Hence, $W(A)$ has a vertex at λ , and often has the appearance of an ice cream cone, as illustrated by the examples reported in [11, Fig. 4]. If, in addition,

$$|p(\lambda)| > |p(v)| \text{ for all } v \in W(A), v \neq \lambda,$$

so that that $Z(p, A)$ consists only of the single point λ , and

$$|p(\lambda)| > \|p(B)\|_2,$$

it is immediate that both the numerator and denominator of the Crouzeix ratio are $|p(\lambda)|$, so $f(p, A) = 1$. Furthermore, we showed in [11, Thm. 2] that f is differentiable at such (p, A) and that its gradient is zero, so that (p, A) is a smooth stationary point (though not necessarily a local minimizer).

4 The computational model

We use the same computational model as in [11], so we briefly explain only the main points. It is well known [15] that $\text{bd } W(A)$, the boundary of $W(A)$, can be characterized as

$$\text{bd } W(A) = \{z_\theta = v_\theta^* A v_\theta : \theta \in [0, 2\pi)\} \tag{6}$$

where v_θ is a normalized eigenvector corresponding to the largest eigenvalue of the Hermitian matrix

$$H_\theta = \frac{1}{2} (e^{i\theta} A + e^{-i\theta} A^*).$$

The proof uses a supporting hyperplane argument [10, Prop. 2]. To accurately and efficiently approximate $\text{bd } W(A)$, we use Chebfun [9], a system for approximating functions on a real interval to machine precision accuracy by adaptive Chebyshev approximation. Chebfun’s function `fov` computes a complex-valued “chebfun”

approximating the extreme points of $\text{bd } W(A)$ on the interval $[0, 2\pi]$, generating interpolation points $\theta \in (0, 2\pi)$ automatically. A second output argument returns any line segments in the boundary as well.²

For the optimization calculations, we use the BFGS method, devised independently in 1970 by Broyden, Fletcher, Goldfarb and Shanno for unconstrained optimization of differentiable functions, but which is also extremely effective for nonsmooth optimization [16]. BFGS requires the computation of the Crouzeix ratio $f(p, A)$ and its gradient at a sequence of iterates generated by the method. The main cost in computing $f(p, A)$ is that of constructing the chebfun representing $\text{bd } W(A)$. Computing the numerator of (2) is then done by invoking two MATLAB functions that have been overloaded to be applicable to chebfuns, namely `polyval` and `norm(. , inf)`, while computation of the denominator, the 2-norm of $p(A)$, is carried out by calls to two standard MATLAB functions, `polyvalm` and `norm`. Once $f(p, A)$ has been computed, the additional computation required to obtain its gradient is minimal, even though the formula is complicated: see [11, Theorem 1] for details. In order to compute the gradient, we need to know $Z(p, A)$, which tells us where $\|p\|_{W(A)}$ is attained; this information is returned by Chebfun's `norm(. , inf)` function. A natural question is: what is the method to do if $Z(p, A)$ contains multiple points, and hence f is not differentiable at (p, A) ? The answer is that since BFGS uses a “gradient paradigm”, not a “subgradient paradigm” [1, 16], this possibility, which is essentially impossible to check exactly in finite precision, is simply ignored. In practice, the algorithm will virtually never compute pairs (p, A) where $Z(p, A)$ contains multiple points, except in the limit. Clearly, small changes in (p, A) may result in large changes in the computed gradient, but this is inherent in nonsmooth optimization, and in fact explains to a large extent why BFGS works so well in this context [16, p. 130]. The BFGS method is a line-search descent method, meaning that at every iteration it uses an inexact line search to repeatedly evaluate the minimization objective f along a search direction in the variable space until the so-called Armijo-Wolfe conditions are satisfied. If, due to rounding errors, it is not possible to satisfy these conditions in a reasonable number of steps, the BFGS method is terminated.

5 Verifying approximate nonsmooth stationarity

In order to verify approximate nonsmooth stationarity of computed pairs (p, A) *a posteriori*, we need to consider some sort of approximate subdifferential of f at (p, A) . It is well known [17, Thm. 10.31] that the subdifferential of a max function on an interval is the convex hull of gradients of the component functions evaluated at points where the max is attained. Hence the importance of the set $Z(p, A)$, which underlies the derivation of the subgradients of the Crouzeix ratio at the Crabb matrix configuration (\tilde{p}, \tilde{A}) given in [10]. We now introduce what may be a novel idea for approximating the subdifferential of a max function on an interval: instead of the convex hull of the gradients evaluated where the max is attained exactly, we

² This modification to `fov` was written by the author.

use the convex hull of gradients evaluated at *local* maximizers for which the locally maximal value is sufficiently close to the globally maximal value. For $\epsilon \geq 0$, define

$$Z_\epsilon(p, A) = \{z \in \text{bd } W(A) : z \text{ is a local maximizer of } \zeta \mapsto |p(\zeta)| \text{ and } |p(z)| \geq (1 - \epsilon)\|p\|_{W(A)}\}. \tag{7}$$

Clearly, $Z_0(p, A) = Z(p, A)$. Then, replace $Z(p, A)$ in [10, Eq. (9)] by $Z_\epsilon(p, A)$ and use the resulting modified formula for $\partial f(p, A)$ in [10, Thm. 3] as our approximate subdifferential, say $\partial^\epsilon f(p, A)$. This should not be confused with other usages of the word ‘‘approximate’’ in subdifferential analysis which sometimes mean ‘‘limiting’’ [17, p. 347], or by approximating the subdifferential using perturbations to (p, A) [2, 3, 12].

Given a computed pair (p, A) , we use Chebfun’s $\text{max}(\cdot, \text{'local'})$ function to compute all local maximizers of $|p|$ on $\text{bd } W(A)$. If $|p|$ is constant on $\text{bd } W(A)$, which can only happen if $W(A)$ is a disk, as in the case that A is a Crabb matrix, Chebfun returns no local maximizers, so we forgo computing the stationarity measure in this case. At all other local minimizers of f such as those described below, the number of local maximizers of $|p|$ on $\text{bd } W(A)$ is necessarily finite, and hence the set $Z_\epsilon(p, A)$ is a discrete set, as opposed to a continuum that would be obtained if we eliminate the ‘‘local maximizer’’ condition in (7). Note also that if $W(A)$ has the ice-cream-cone configuration, with $|p|$ maximized on $\text{bd } W(A)$ only at the vertex λ , then the local maximizers returned by Chebfun do not include λ , since $\text{bd } W(A)$ is not smooth there. Hence, it is important to compute the global maximum using Chebfun’s $\text{max}(\cdot)$ as well as the local maximizers with $\text{max}(\cdot, \text{'local'})$.

Finally, since the exact nonsmooth stationarity condition is $0 \in \partial f(p, A)$, we compute

$$d = \text{argmin}\{\|g\|_2 : g \in \partial^\epsilon f(p, A)\}, \tag{8}$$

the solution of a convex quadratic programming problem, and use $\|d\|_2$ as a measure of approximate nonsmooth stationarity.

6 The new computational results

M. Hairer³ pointed out that by starting with initial data generated from the normal distribution, as done in the results reported in [11], we might be biasing the optimization towards matrices whose fields of values are disks. He suggested trying initial data generated from distributions with heavy tails, e.g., numbers of the form

$$xe^{\alpha x^2} \tag{9}$$

where x is obtained from the standard normal distribution and $\alpha > 0$. Indeed, experiments show that for fixed n , as α increases, it becomes much less likely that BFGS will generate Crabb matrix configurations with fields of values a disk and Crouzeix

³ Private communication, 2019

ratio 0.5. But instead of finding lower values that would disprove the conjecture, what happens instead is that BFGS is much more likely to find other locally minimal values of the Crouzeix ratio between 0.5 and 1. In the experiments reported here, we fix $\alpha = 2$ in (9). This is large enough to discover many more locally minimal values than using $\alpha = 0$ or $\alpha = 1$, but not so large that we have problems with overflow.

In the results reported here, in addition to optimizing over real polynomials and real matrices, we also present results obtained by optimizing over complex polynomials and complex matrices. Without loss of generality, since the Crouzeix ratio is invariant to unitary similarity transformations of the matrix, in the real case we restrict the matrices to upper Hessenberg form; in the complex case, we restrict them to upper triangular form. Furthermore, in the real case, since the field of values is symmetric with respect to the real axis, we compute its boundary only in the closed upper half-plane, restricting θ in (6) to the interval $[\pi, 2\pi]$, and modifying the definitions of $Z(p, A)$ and $Z_c(p, A)$ to restrict them to points z with $\text{Re}(z) \geq 0$.

6.1 The case $n = 2$

Figure 1 shows the final values of the Crouzeix ratio f for $n = 2$ with the maximal degree m ranging over 2, 3, 4 and 5, obtained by running BFGS⁴ initialized from 10,000 randomly generated starting points using (9) with $\alpha = 2$, for each of the real and complex cases. The blue dots show the *sorted* final values of the Crouzeix ratio obtained from optimizing over real matrices of order n and real polynomials with maximal degree m , and the red dots show the *sorted* final values when optimizing over complex matrices of order n and complex polynomials with maximal degree m . Note that since the results for the real case are plotted first, many of the blue dots are overwritten by red dots. The plateaus clearly indicate locally minimal values, as these values are found repeatedly from many starting points. Furthermore, we see that most of the locally minimal values found are the *same* for both the real and complex cases. That said, there is not a great deal of similarity between the real and complex results. The width of the plateaus of locally minimal values found represents the frequency with which they are found, and this varies greatly between the real and complex cases. Partly for this reason, it is difficult to check whether there is a one-to-one correspondence between the locally minimal values found in the real and complex cases; there are some possible counterexamples if we assume that 10,000 runs is enough to see these features, but obviously we have no basis for this. The locally minimal values 0.5 and 1 are clearly apparent in every case. The former, 0.5, is known to be globally minimal for $n = 2$, and is globally minimal for all n and m if Crouzeix's conjecture is true. The latter value, 1, is the stationary value associated with ice-cream-cone configurations discussed above.

It's interesting to see how the other locally minimal values visible in Fig. 1 vary with m . In both the real and complex cases, in the case $m = 3$, the widest plateau between 0.5 and 1 has value 0.84375, and this is also visible as shorter

⁴ Using HANSO 2.2 (www.cs.nyu.edu/overton/hanso), with `options.normtol = 1e-8`.

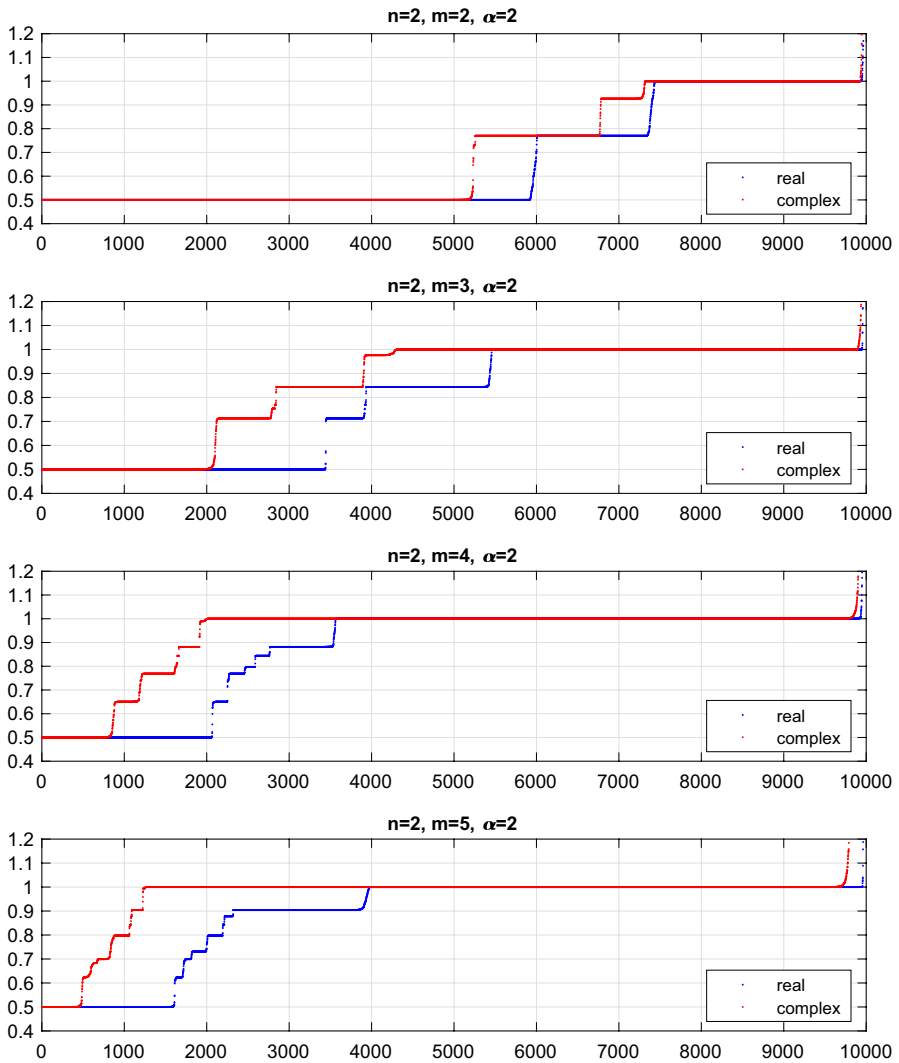


Fig. 1 Sorted final values of the Crouzeix ratio f obtained from 10,000 runs for $n = 2$, $m = 2, 3, 4, 5$, optimizing over *real* $n \times n$ matrices and *real* polynomials of degree at most m (blue dots) and optimizing over *complex* $n \times n$ matrices and *complex* polynomials of degree at most m (red dots), with starting points generated by the heavy-tail distribution defined by (9), with $\alpha = 2$. Note that many of the blue dots are overwritten by red dots

plateaus in the cases $m = 4$ and $m = 5$, but it does not appear in the case $m = 2$. Meanwhile, again in both the real and complex cases, in the case $m = 2$, the widest plateau between 0.5 and 1, with value 0.7698, is reduced to much narrower plateaus in the case $m = 3$, $m = 4$ and $m = 5$, which are only visible with MATLAB’s zoom tool. This behavior is unexpectedly complicated.

Table 1 Four locally minimal values for $n = 2$, $m = 3$, including 0.5 and 1

	Run #	Numer	Denom	f	$ Z_\epsilon $	$\ d\ $
R	1	4.195e+03	8.391e+03	0.5000000000		
R	3400	1.031e+05	2.061e+05	0.5000188514		
R	3500	7.808e+06	1.095e+07	0.7132185867	1	3.313e-08
R	3850	2.371e+11	3.325e+11	0.7132189899	1	4.689e-09
R	4000	7.904e+11	9.368e+11	0.8437496323	1	8.131e-09
R	5200	9.164e+17	1.086e+18	0.8437501759	1	5.790e-09
R	5500	5.624e+00	5.624e+00	1.0000000000	1	8.436e-16
R	9800	6.694e+14	6.694e+14	1.0000000084	1	4.322e-09
C	1	1.362e+15	2.724e+15	0.5000000000		
C	1900	5.001e+03	9.996e+03	0.5002960578		
C	2200	3.113e+07	4.365e+07	0.7132064490	2	1.574e-05
C	2700	1.122e+10	1.574e+10	0.7132194699	2	1.263e-07
C	2900	7.691e+08	9.115e+08	0.8437482418	2	1.809e-07
C	3800	2.371e+11	2.810e+11	0.8437500721	2	5.533e-09
C	4500	2.470e+07	2.470e+07	1.0000000000	3	1.845e-04
C	9800	7.825e+18	7.825e+18	1.0000004701	1	3.734e-09

The first column indicates whether the data is from the real or complex run, and the second column shows the relevant run number. The next columns show the numerator and denominator of the final Crouzeix ratio, as well as the ratio f itself. The final two columns show the number of points in $Z_\epsilon(p, A)$ and the resulting approximate stationarity measure, using $\epsilon = 10^{-4}$.

Table 1 shows more details associated with four locally minimal values clearly visible as plateaus in the second panel of Fig. 1, for $n = 2$, $m = 3$, namely, 0.5, 0.713, 0.844 and 1. These locally minimal values are found in both the real case (indicated by R in the first column of the table) and the complex case (indicated by C in the table).⁵ The second column of the table indicates the run number in the horizontal scale used in Fig. 1. These indices are chosen to *roughly* correspond to the first and last points in each plateau. Consequently, corresponding pairs of values of the Crouzeix ratio f , shown in the fifth column, indicate the approximate precision to which the locally minimal values are found. The third and fourth columns show the associated numerator and denominator of (2). Note that, despite the enormous values of the numerator and denominator of the Crouzeix ratio, the ratio f is consistently computed to several digits of agreement even over large numbers of starting points. For example, the first 3400 real runs approximate 0.5 to 4 digits of accuracy, while the first 1900 complex runs approximate 0.5 to 3 digits. The final two columns of Table 1 show two quantities associated with the approximate stationarity measure described in Sect. 5, namely, the number of points in $Z_\epsilon(p, A)$ and the 2-norm of the corresponding vector obtained in (8), using $\epsilon = 10^{-4}$. For the runs which find f approximately equal to 0.5, these values are omitted, because, at the Crabb matrix

⁵ The complex runs also identify a fifth locally minimal value, 0.977, but it is not clear whether this is a locally minimal value in the real case; in any case, we cannot conclude that from the figure.

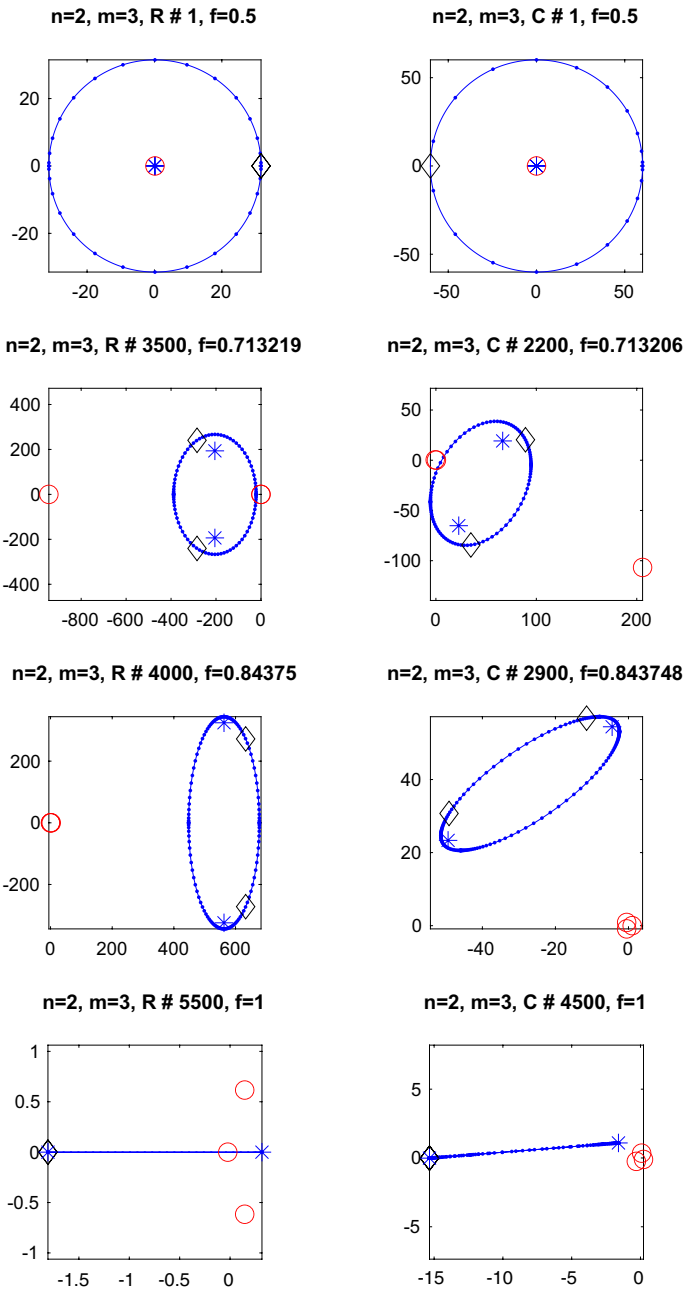


Fig. 2 Configurations of some local minimizers obtained for $n = 2$ and $m = 3$ with locally minimal values 0.5, 0.713, 0.844 and 1, corresponding to real runs 1, 3500, 4000 and 5500 on the left and complex runs 1, 2200, 2900 and 4500 on the right. Solid blue curves show the boundaries of the field of values $W(A)$, blue asterisks show the eigenvalues of A , red circles show the roots of p , and black diamonds show the points on $W(A)$ where $\|p\|_{W(A)}$ is attained

configurations, $Z(p, A)$ is a continuum, and, as noted earlier, Chebfun does not find any local maxima of $|p|$ on $\text{bd } W(A)$ in this case. In all other cases, the value of $\|d\|$ is small, strongly indicating that the computed (p, A) is an approximate stationary point. We comment further below on the number of points in $Z_e(p, A)$.

Figure 2 shows the fields of values of A for the final configuration obtained by the first real run (left) and first complex run (right) given in Table 1 associated with the locally minimal values 0.5, 0.713, 0.844 and 1 (top row, second row, third row and bottom, respectively). In the real case, the fields of values are necessarily symmetric w.r.t. the real axis. Blue asterisks denote the eigenvalues of A , red circles show the roots of p , and black diamonds show the points in $Z_e(p, A)$, where $\|p\|_{W(A)}$ is (nearly) attained. Clearly, when shifted, scaled and rotated, the final fields of values found by the complex runs are very similar to the ones found by the real runs that approximate the same locally minimal value.

The known configuration attaining $f = 0.5$ is (p, A) where A is a Jordan block (the Crabb matrix for $n = 2$), for which $W(A)$ is a disk, and $p(z) = z - \lambda$, where λ is the double eigenvalue of A . We see this configuration in the top two panels of Fig. 2. Note that the two eigenvalues of the final computed matrix A and one of the roots of the final computed polynomial p (nearly) coincide at the center of the disk. Since the computations are being done with maximal degree $m = 3$, there are two other roots of p , which, for the optimal p , are ∞ , but, in our computations, are numbers with very large modulus that are not shown. In theory, the optimal $\|p\|_{W(A)}$ is attained at every point on the boundary of $W(A)$, but only one point is shown.

In the second row of Fig. 2, the eigenvalues of A are distinct, so $W(A)$ is elliptical, but two of the three roots of p (nearly) coincide, a little outside $W(A)$, while the third root is on the other side of and further away from $W(A)$. In the third row, $W(A)$ is a more eccentric ellipse than in the second row, and the three roots of p are all (nearly) coincident.⁶ These are interesting, and decidedly non-random, configurations. In the second and third rows, $\|p\|_{W(A)}$ is attained at two points on the boundary of $W(A)$. In the real cases, this is a consequence of the imposed structure: since $\|p\|_{W(A)}$ is attained at a complex point, it must also be attained at the conjugate point. For this reason, as shown in Table 1, $Z_e(p, A)$ contains only one point, since only points in the closed upper half-plane are admissible in the real case. It follows that the minimizer is a smooth stationary point in the real data space, and the vector d whose norm is shown in the table is actually the gradient. On the other hand, in the complex case, no such structure is imposed, and the double attainment is reflected by $Z_e(p, A)$ having two points in this case. Hence, this is a nonsmooth minimizer in the complex data space, and the vector d whose norm is shown in Table 1 is an approximate subgradient, not a gradient.

The final configuration for the locally optimal value 1 is a diagonal matrix whose field of values is a line segment, which, as $n = 2$, is a degenerate case of the ice-cream-cone fields of values mentioned earlier. In the bottom row of Fig. 2, in the

⁶ Since optimizing over complex matrices and complex polynomials requires much more computation than in the real case, the results are likely less accurate in the complex case, as is suggested by the three roots of p being less close to coincident.

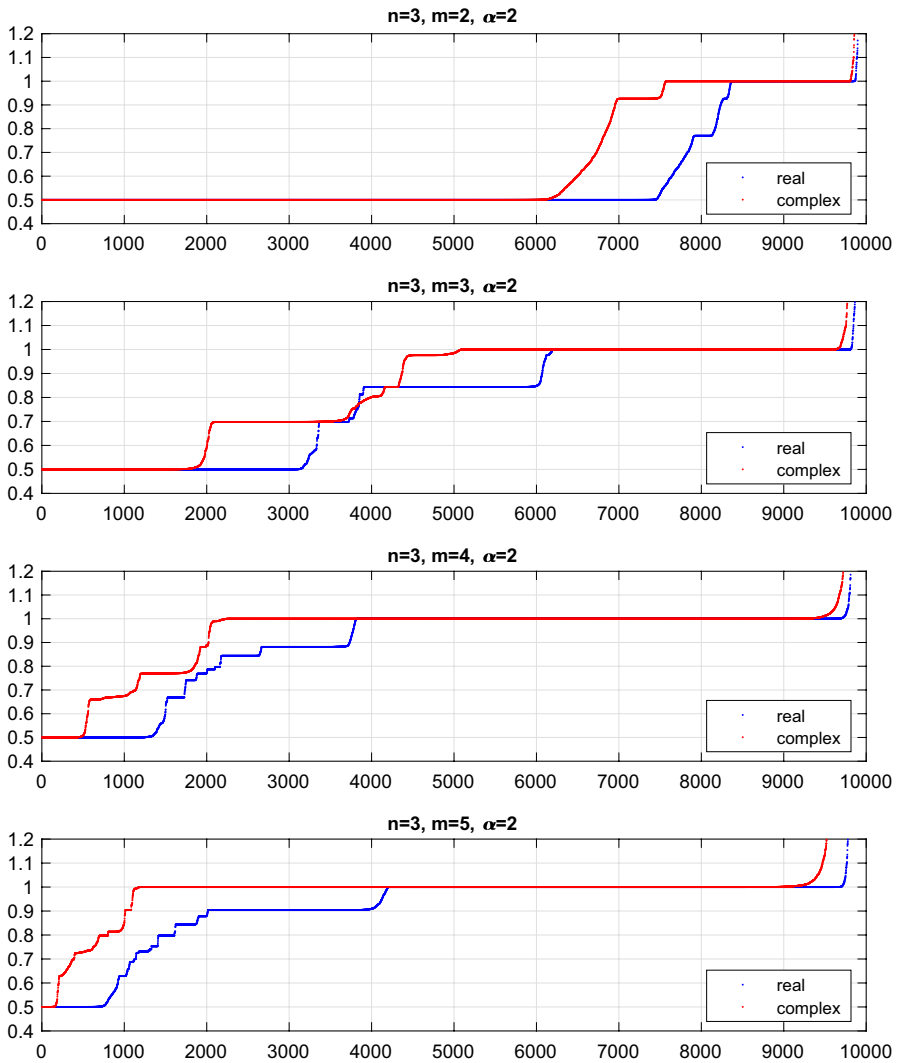


Fig. 3 Sorted final values of the Crouzeix ratio f obtained from 10,000 runs for $n = 3, m = 2, 3, 4, 5$, optimizing over *real* $n \times n$ matrices and *real* polynomials of degree at most m (blue dots) and optimizing over *complex* $n \times n$ matrices and *complex* polynomials of degree at most m (red dots), with starting points generated by the heavy-tail distribution defined by (9), with $\alpha = 2$. Note that many of the blue dots are overwritten by red dots

real case, indeed we see that $W(A)$ is a line segment, so $Z_\epsilon(p, A)$ has a single point, f is differentiable and the vector d is a gradient, with small norm. On the other hand, although the field of values computed in the complex case is nearly a line segment, it is not exactly a line segment, and $Z_\epsilon(p, A)$ actually has 3 points, although they are all nearly identical. Consequently, the vector d is an approximate subgradient, but it still has small norm, showing the robustness of the calculations.

Table 2 Four locally minimal values for $n = 3$, $m = 3$, including 0.5 and 1

	Run #	Numer	Denom	f	$ Z_c $	$\ d\ $
R	1	1.607e+14	3.213e+14	0.5000000000		
R	3000	4.982e+02	9.962e+02	0.5001593709		
R	3400	5.766e+09	8.264e+09	0.6978015654	2	6.937e-04
R	3700	3.127e+17	4.481e+17	0.6978024061	2	1.022e-05
R	4000	2.395e+12	2.839e+12	0.8437498418	1	1.370e-08
R	5800	6.651e+16	7.883e+16	0.8437562126	2	3.681e-08
R	6250	2.489e+03	2.489e+03	1.0000000000	1	6.227e-12
R	9800	7.920e+19	7.920e+19	1.0000670371	1	3.732e-09
C	1	4.262e+13	8.523e+13	0.5000000000		
C	1400	5.786e+10	1.157e+11	0.5001167444		
C	2200	7.823e+10	1.121e+11	0.6978004851	3	6.705e-05
C	3200	4.851e+19	6.950e+19	0.6979798248	3	9.135e-07
C	4200	4.439e+16	5.261e+16	0.8437493557	2	1.055e-08
C	4300	6.078e+19	7.204e+19	0.8437505791	2	7.263e-09
C	5500	1.005e+14	1.005e+14	1.0000000000	1	3.356e-13
C	9500	3.181e+19	3.181e+19	1.0000027774	1	5.767e-09

The first column indicates whether the data is from the real or complex run, and the second column shows the relevant run number. The next columns show the numerator and denominator of the final Crouzeix ratio, as well as the ratio f itself. The final two columns show the number of points in $Z_c(p, A)$ and the resulting approximate stationarity measure, using $\epsilon = 10^{-4}$.

6.2 The case $n = 3$

Figure 3 shows sorted final values for the case $n = 3$, for $m = 2, 3, 4, 5$, again using 10,000 starting points for each of the real and complex runs. Let us again focus on the results for $m = 3$, shown in the second panel. Table 2 shows details associated with the four locally minimal values 0.5, 0.698, 0.844 and 1 that are clearly visible in the figure.⁷ The locally minimal value 0.698 is significantly less than the value 0.713 observed in Fig. 1 for the case $n = 2$, $m = 3$, but the locally minimal value 0.844 is the same as the value observed in the case $n = 2$, $m = 3$. Figure 4 shows the fields of values of A for the final configuration obtained by the first real run (left) and first complex run (right) given in Table 2 associated with the four locally minimal values 0.5, 0.698, 0.844 and 1 (top row, second row, third row and bottom, respectively).

The top two panels again show that $W(A)$ is a disk. In the top left (the real case), the three eigenvalues of A are (nearly) coincident with two of the roots of p , approximating a configuration (p, A) where A is a Crabb matrix of order 3 (a 3×3 Jordan block) and $p(z) = (z - \lambda)^2$, where λ is the eigenvalue of A . Zooming in indeed shows

⁷ Again, we omit a fifth locally minimal value, 0.977, that is clearly identified by the complex runs; unlike in the case $n = 2$, $m = 3$, zooming in indicates that there is also a small plateau with this value for the real runs.

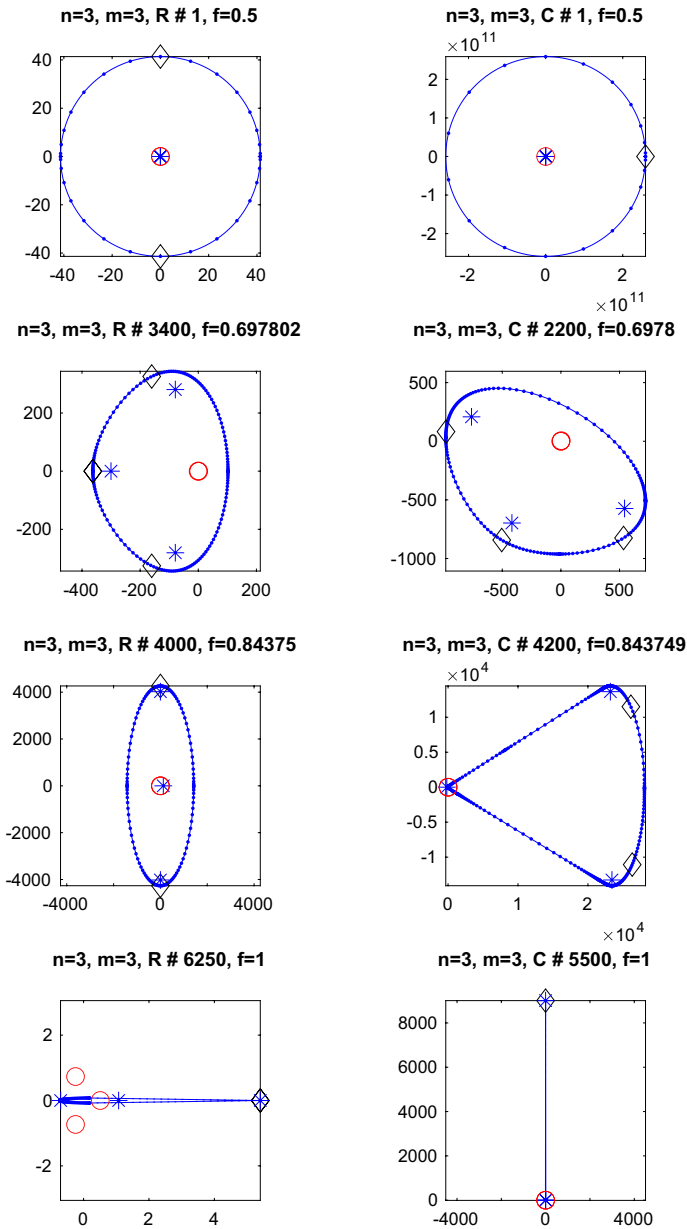


Fig. 4 Configurations of some local minimizers obtained for $n = 3$ and $m = 3$ with locally minimal values 0.5, 0.698, 0.844 and 1, corresponding to real runs 1, 3400, 4000 and 6250 on the left and complex runs 1, 2200, 4200 and 5500 on the right. Solid blue curves show the boundaries of the field of values $W(A)$, blue asterisks show the eigenvalues of A , red circles show the roots of p , and black diamonds show the points on $W(A)$ where $\|p\|_{W(A)}$ is attained

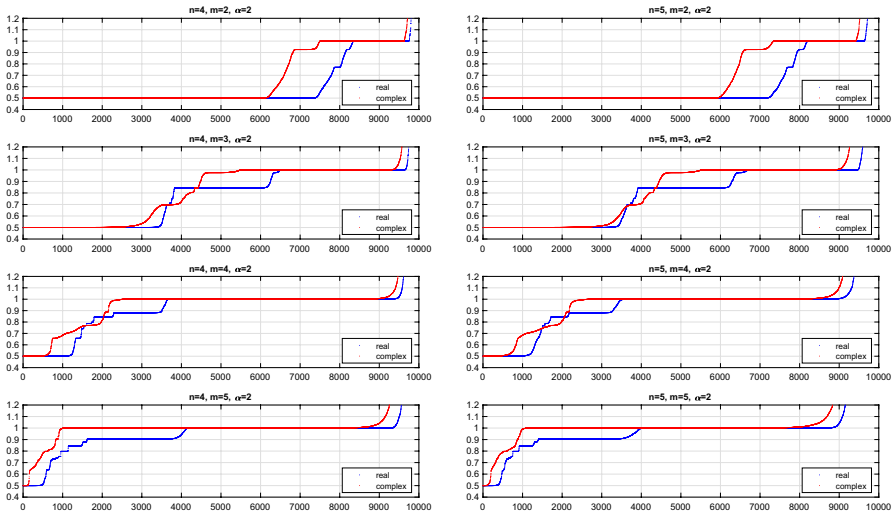


Fig. 5 Sorted final values of the Crouzeix ratio f obtained for $n = 4, 5, m = 2, 3, 4, 5$

that the three computed eigenvalues of A and two of the roots of p are nearly coincident, while the third root (not shown) has enormous modulus. However, in the top right (the complex case), the optimal configuration that is approximated is subtly different. The computed complex triangular matrix A has one upper triangular entry which is much larger than the others, and only one of the roots of the computed p is close to the eigenvalues of A , while the other two (not shown) have very large modulus. This indicates that the optimal configuration being approximated is (p, A) where A is block diagonal with a 2×2 Jordan block J (the Crabb matrix of order 2) and $p(z) = z - \lambda$, where λ is the double eigenvalue of J , with the third eigenvalue of A separated from the others but inside the field of values of J . Indeed, zooming in we find that one root of p and two eigenvalues of A are nearly coincident, with the third eigenvalue separated from them, though not by much.

In the second row of Fig. 4, we see a new configuration: the three eigenvalues of A are well separated but the three roots of p are (nearly) coincident, and $\|p\|_{W(A)}$ is (nearly) attained at three points on the boundary. Consequently, as we see in Table 2, in the real case Z_ϵ has two points (one real and one complex) while in the complex case Z_ϵ contains three complex points. Again, we see from Table 2 that $\|d\|$ is small in both cases, indicating approximate nonsmooth stationarity.

In the third row, associated with the locally minimal value 0.844 that was observed earlier, the fields of values clearly indicate that the associated matrices are nearly block diagonal, with blocks of order 2 and 1. On the left, the eigenvalue corresponding to the 1×1 block lies inside the field of values of the 2×2 block, while on the right, it does not; however, in both cases the numerator and denominator of the Crouzeix ratio are determined by the (approximate) 2×2 block; the other block is “inactive”. This is the reason why the locally minimal value associated with the two lower panels is the same as that found for the case $n = 2, m = 2$. In contrast, the fields of values shown in the second row are clearly *not* associated with block

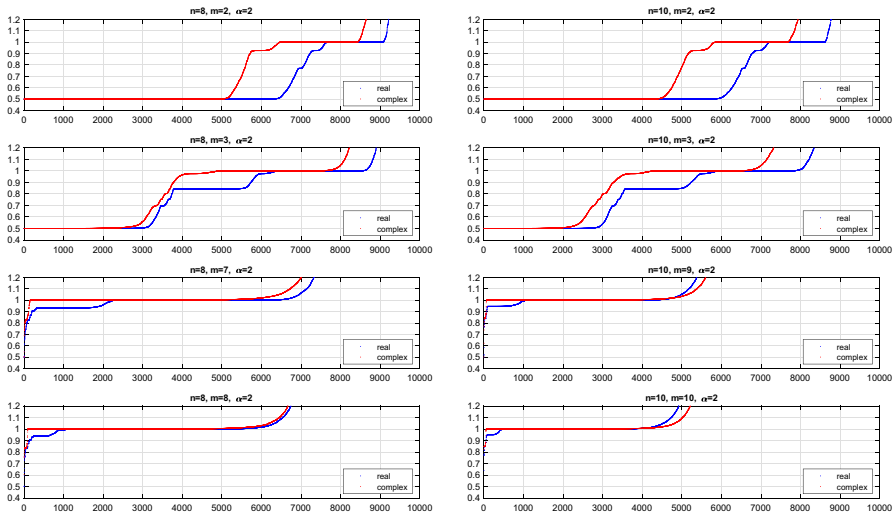


Fig. 6 Sorted final values of the Crouzeix ratio f obtained for (left) $n = 8, m = 2, 3, 7$ and 8 , and (right) $n=10, m = 2, 3, 9$ and 10

diagonal matrices: thus, the associated locally minimal value is *not* a locally minimal value when $n = 2$. To emphasize this point, we might say that 0.698 is a *genuine* locally minimal value for $n = 3$.

Finally, in the fourth row we see two very elongated ice-cream-cone configurations, which are smooth stationary points.

6.3 Cases with $n \geq 4$

Figure 5 shows results for $n = 4$ (left) and $n = 5$ (right), with $m = 2, 3, 4, 5$. Again we see strong evidence of locally minimal values between 0.5 and 1 , with at least some values coinciding for the real and complex cases. Figure 6 shows results for $n = 8, m = 2, 3, 7, 8$ (left), and $n=10, m = 2, 3, 9, 10$ (right). For these values of n , we can still observe locally minimal values between 0.5 and 1 when $m = 2$ or 3 , but for larger m only a handful of results with f below 1 are observed.

7 Concluding remarks

Crouzeix’s conjecture [7] states that the globally minimal value of the Crouzeix ratio (2) is 0.5 , regardless of n and m , and it was demonstrated in [11] that 1 is a frequently occurring locally minimal value. Making use of the heavy-tailed distribution (9) to initialize our optimization runs, we have demonstrated for the first time that the Crouzeix ratio has many other locally minimal values between 0.5 and 1 , even for $n = 2, m = 3$ and $n = 3, m = 3$, cases that we studied in detail. Not only did we show that the same function values are repeatedly obtained for many different starting points, but we also verified that approximate nonsmooth stationarity conditions

hold at computed candidate local minimizers. We also found that the same locally minimal values are often obtained both when optimizing over real matrices and polynomials, and over complex matrices and polynomials. The appearance of so many nonsmooth local minimizers suggests how very complex a function the Crouzeix ratio is and could perhaps shed some light on why mathematically establishing its global minimum is so difficult.

We think that minimization of the Crouzeix ratio makes a very interesting nonsmooth optimization case study illustrating among other things how effective the BFGS method is for nonsmooth optimization. Our method for verifying approximate nonsmooth stationarity is based on what may be a novel approach to finding approximate subgradients of max functions on an interval, exploiting Chebfun's ability to efficiently find local maximizers on intervals.

Our extensive computations strongly support Crouzeix's conjecture. We have presented results for nearly half a million optimization runs reported in Figs. 1, 3, 5 and 6, computing the Crouzeix ratio for about 250 million pairs (p, A) . The computations were done using the high performance computing cluster at New York University, running the code on hundreds of CPU cores using a `parfor` "parallel for" loop in `MATLAB`. Doing this, even the 80,000 runs for $n = 10$ took less than 3 hours. We always found that the smallest locally minimal value was 0.5.

References

1. Asl, A., Overton, M.L.: Behavior of limited memory BFGS when applied to nonsmooth functions and their Nesterov smoothings. In: Al-Baali, M., Grandinetti, L., Purnama, A. (eds.) *Recent Dev. Numer. Anal. Optim.* Springer, Cham (2021)
2. Burke, J.V., Lewis, A.S., Overton, M.L.: Approximating subdifferentials by random sampling of gradients. *Math. Oper. Res.* **27**, 567–584 (2002)
3. Burke, J.V., Lewis, A.S., Overton, M.L.: A robust gradient sampling algorithm for nonsmooth, non-convex optimization. *SIAM J. Optim.* **15**, 751–779 (2005)
4. Clarke, Frank H.: Generalized gradients and applications. *Trans. Amer. Math. Soc.* **205**, 247–262 (1975)
5. Crouzeix, M., Palencia, C.: The numerical range is a $(1 + \sqrt{2})$ -spectral set. *SIAM J. Matrix Anal. Appl.* **38**(2), 649–655 (2017)
6. Crabb, Michael J.: The powers of an operator of numerical radius one. *Michigan Math. J.* **18**, 253–256 (1971)
7. Crouzeix, M.: Bounds for analytical functions of matrices. *Integral Eq. Oper. Theory* **48**, 461–477 (2004)
8. Crouzeix, Michel: Some constants related to numerical ranges. *SIAM J. Matrix Anal. Appl.* **37**(1), 420–442 (2016)
9. Driscoll, T.A., Hale, N., Trefethen, L.N.: *Chebfun guide*. Pafnuty Publications, Oxford (2014)
10. Greenbaum, A., Lewis, A.S., Overton, M.L.: Variational analysis of the Crouzeix ratio. *Math. Program. Ser. A* **164**(1–2), 229–243 (2017)
11. Greenbaum, A., Overton, M.L.: Numerical investigation of Crouzeix's conjecture. *Linear Algebra Appl.* **542**, 225–245 (2018)
12. Goldstein, A.A.: Optimization of Lipschitz continuous functions. *Math. Progr.* **13**, 14–22 (1977)
13. Guglielmi, N., Overton, M.L., Stewart, G.W.: An efficient algorithm for computing the generalized null space decomposition. *SIAM J. Matrix Anal. Appl.* **36**(1), 38–54 (2015)
14. Horn, R.A., Johnson, C.R.: *Topics in matrix analysis*. Cambridge University Press, Cambridge, U.K. (1991)

15. R. Kippenhahn. Über den Wertevorrat einer Matrix. *Math. Nachr.*, 6:193–228, 1951. English translation by P.F. Zachlin and M.E. Hochstenbach, *Linear and Multilinear Algebra* 56, pp. 185–225, 2008
16. Lewis, A.S., Overton, M.L.: Nonsmooth optimization via quasi-Newton methods. *Math. Program. Ser. A* **141**(1–2), 135–163 (2013)
17. R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*, volume 317 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1998

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.