Partial smoothness of the numerical radius at matrices whose fields of values are disks

A.S. Lewis[*] M.L. Overton[†]

December 31, 2018

Abstract

Solutions to optimization problems involving the numerical radius often belong to a special class: the set of matrices having field of values a disk centered at the origin. After illustrating this phenomenon with some examples, we illuminate it by studying matrices around which this set of "disk matrices" is a manifold with respect to which the numerical radius is partly smooth. We then apply our results to matrices whose nonzeros consist of a single superdiagonal, such as Jordan blocks and the Crabb matrix related to a well-known conjecture of Crouzeix. Finally, we consider arbitrary complex three-by-three matrices; even in this case, the details are surprisingly intricate. One of our results is that in this real vector space with dimension 18, the set of disk matrices is a semi-algebraic manifold with dimension 12.

1 Introduction: the numerical radius

Given a matrix A in the space $\mathbf{M}^n$ of $n \times n$ complex matrices, the asymptotic convergence of the discrete-time dynamical system

$$x_{k+1} = Ax_k \qquad (k = 0, 1, 2, \ldots)$$

is governed by the spectrum of A. Specifically, the iterates $x_k \in \mathbf{C}^n$ converge to zero for all initial points x_0 if and only if A is *stable*, meaning its spectrum is contained in the open unit disk. However, to understand the *transient* behavior of the iterates

— how $\|x_k\|$ depends on k — the spectrum is inadequate. Instead we can rely on the following well-known inequality:

$$(1.1) \qquad \|A^k\|_2 \le 2\big(r(A)\big)^k \qquad (k = 0, 1, 2, \ldots),$$

where $\| \cdot \|_2$ denotes the spectral norm (the largest singular value), $r(A)$ denotes the numerical radius of A,

$$r(A) = \max\{|z| : z \in W(A)\},$$

and the set

$$W(A) = \{u^* A u : u \in \mathbf{C}^n, \|u\|_2 = 1\}$$

is the field of values (numerical range) of A. For an excellent discussion of this and related issues, see Trefethen and Embree's book [34, §14,17].

The numerical radius is a vector space norm, but it does not satisfy the submultiplicative property required in most definitions of a matrix norm [20]. However, Lax and Wendroff, in their famous 1964 paper on difference equations for hyperbolic PDEs [23], showed that if $r(A) \le 1$, then the sequence (A^k) is bounded. Halmos then conjectured the "power inequality"

$$r(A^k) \le \big(r(A)\big)^k \qquad (k = 0, 1, 2, \ldots).$$

The first proof, by Berger [3], was simplified by Pearcy [4,15,22,31]. The inequality (1.1) is an elementary consequence [20, p. 379]. A variety of other properties of the numerical radius are collected in [27, §18.4], and a useful bibliography appears in [26]. Research on the topic remains very active.[1]

An elegant proof of the power inequality relies on the following semidefinite representation of the numerical radius due to [29], based in part on [2]. We denote the space of n-by-n Hermitian matrices by $\mathbf{H}^n$.

Theorem 1.2 (Semidefinite representability of numerical radius)
The numerical radius of a matrix $A \in \mathbf{M}^n$ is bounded above by $\alpha \in \mathbf{R}$ if and only if there exists a matrix $Z \in \mathbf{H}^n$ such that the Hermitian matrix

$$\begin{bmatrix} \alpha I + Z & A \\ A^* & \alpha I - Z \end{bmatrix}$$

is positive semidefinite.

---

[1]So far this century, at least 150 papers have appeared with titles including the string "numerical radius" (according to MathSciNet). The concept of numerical radius goes back at least to Lumer (1961) [28], but may well be older.

This representation in particular allows us to compute the numerical radius of A by semidefinite programming, choosing variables α and Z to minimize α subject to the semidefinite constraint. For more direct methods, see [30, 36].

The numerical radius appears naturally in matrix optimization and best approximation problems: one survey, from an operator-theoretic perspective, appears in [1]. Our current work is motivated by the following observation: optimization involving the numerical radius quite commonly results in matrices whose fields of values are disks (by which we mean disks in the complex plane centered at zero, including the trivial case $\{0\}$). We call such matrices *disk matrices*, and denote the set of disk matrices by $\mathcal{A}$. Surprisingly, an explicit parametrization of disk matrices was not known until recent work of Crouzeix [10], discussed below in Section 8, and based on much older work of Ando [2].

Disk matrices are relatively rare: we focus here primarily on examples when, locally, the set $\mathcal{A}$ has the structure of a manifold of codimension $2n$ in the $2n^2$-dimensional real vector space $\mathbf{M}^n$. Our aim here is to illuminate variational properties of the numerical radius (in particular the property of partial smoothness [24]) that, notwithstanding their rarity, produces disk matrices at optimality.

The paper is organized as follows. In Section 2, we discuss two interesting examples where disk matrices appear as solutions to optimization problems involving the numerical radius. The first example involves feedback control, while the second concerns the Crouzeix conjecture in matrix analysis [9, 11]. Then, in Section 3, we introduce the ideas we use from convex analysis, including the notion of partial smoothness. Section 4 summarizes our results for a well-known disk matrix that arises in the study of Crouzeix's conjecture. In Section 5 we characterize subgradients of the numerical radius at disk matrices, and then in Sections 6 and 7 we develop our main results concerning partial smoothness of the numerical radius at disk matrices. Section 8 discusses our results in the framework of the Crouzeix parametrization of disk matrices, and we end in Section 9 with a thorough analysis of the three-by-three case.

2   Examples of numerical radius optimization

For now, we focus on two examples to illustrate our theme. More examples, arising from the proximal operator for the numerical radius, follow in Section 4.

Feedback control

Our first example generalizes our original dynamical system $x_{k+1} = Ax_k$ in the space $\mathbf{C}^n$ by involving an input vector $u_k \in \mathbf{C}^m$ and an output vector $y_k \in \mathbf{C}^p$, related as follows:

$$x_{k+1} = Ax_k + Bu_k, \qquad y_k = Cx_k,$$

for given matrices B and C. If we set $u_k = Ky_k$ for some matrix K (the controller), we arrive at the system

$$x_{k+1} = M(K)x_k \qquad (k = 0, 1, 2, \ldots),$$

defined by the matrix $M(K) = A + BKC$. The classical static output feedback problem [39] is to choose K so that $M(K)$ is stable. The problem of determining whether such a matrix K exists in general seems to be hard; indeed, when bounds are put on all the entries of K, this problem is NP-hard [5]. However, as long as $mp \geq n$ (or $mp > n$ in the real case), the problem *is* generically solvable via "pole placement": K can be chosen to place the eigenvalues of $M(K)$ in any desired location [37, 38].

On the other hand, as we remarked earlier, although stability by itself does not give information about the transient behavior of the system, such information *is* given by the numerical radius $r\big(M(K)\big)$. By contrast with the question of stability, the problem of minimizing the function $r\big(M(\cdot)\big)$ is tractable via semidefinite programming, thanks to Theorem 1.2, observing that the map M is affine.

We now report on some computational experiments minimizing $r\big(M(K)\big)$ which clearly support our claim that minimizing the numerical radius frequently leads to disk matrices. Normally A, B and C depend on the application of interest, but here we generated them randomly, choosing them to be real matrices with i.i.d. entries from the normal distribution. We set $n = 5$ and considered m and p varying from 1 to n. For each choice of (m, p), we generated 1000 random triples (A, B, C) with dimensions $n \times n$, $n \times m$ and $p \times n$ respectively. We then used the CVX package [16] (with precision set to "best") to minimize the numerical radius $r\big(M(\cdot)\big)$ over m-by-p matrices K for each randomly generated triple, resulting in an optimal controller $\bar{K}$. We then used the Chebfun package [12], as described in [18], to compute the boundary of the field of values of the optimal matrix $\bar{M} = M(\bar{K})$ and measure how close it is to being a circle centered at the origin. To do this we recomputed the numerical radius

$$r(\bar{M}) = \max\big\{|z| : z \in \mathrm{bd}\big(W(\bar{M})\big)\big\}$$

as well as the "inner numerical radius" [7]

$$\tilde{r}(\bar{M}) = \min\big\{|z| : z \in \mathrm{bd}\big(W(\bar{M})\big)\big\},$$

and computed a measure $r(\bar{M}) - \tilde{r}(\bar{M})$ of the distortion of $W(\bar{M})$ from a disk centered at zero. Table 1 shows, for each triple, in how many of the 1000 instances the resulting value $r(\bar{M}) - \tilde{r}(\bar{M})$ was below a threshold that we set to 10^{-7}. These quantities are shown as percentages rounded to the nearest 10% because repeated runs indicated that with this choice of rounding, the results are reasonably consistent. For the cases where 0% is displayed, actually not even one out of the 1000 runs resulted in a matrix satisfying the criterion $r(\bar{M}) - \tilde{r}(\bar{M}) < 10^{-7}$, while for

	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$
$m=1$	0 %	0 %	0 %	0 %	20 %
$m=2$	0 %	0 %	0 %	50 %	40 %
$m=3$	0 %	0 %	70 %	70 %	70 %
$m=4$	0 %	50 %	70 %	90 %	90 %
$m=5$	20 %	40 %	70 %	90 %	100 %

Table 1: Percentage of cases where $W(M(\bar{K}))$ is a disk

the case $m = p = 5$ all 1000 optimal matrices satisfied the given criterion; this is simply because in this case, with 25 free variables, the optimal matrix is zero with probability one.

More interesting is the fact that these computational results are consistent with our theoretical results given later in the paper. In the cases we study, we show that the set of disk matrices in the Euclidean space $\mathbf{M}^n$ of complex $n \times n$ matrices typically has codimension $2n$. (We do not pursue the real case here, but in the space of $n \times n$ real matrices, the analogous codimension is n.) In the cases in the table where none of the optimal matrices satisfied the disk matrix criterion, the number of free variables in the optimization problem, mp, is less than n, and therefore the probability of the affine space of feasible matrices intersecting the manifold of disk matrices is zero. In contrast, in all cases when $n \leq mp < n^2$, one can expect the affine space of feasible matrices to intersect the manifold of disk matrices, so that there is at least the possibility that the optimal matrix is a disk matrix — and evidently, from the results in the table, the probability of this happening is nonzero, ranging in fact from approximately 20% to 90%.

The Crouzeix ratio

Our second example arises from observations [17,18] that minimizing a certain ratio associated with the Crouzeix conjecture [9–11] seems to result in disk matrices. In particular, a special case that is known to be true is that, for matrices $X \in \mathbf{M}^n$, the ratio

$$(2.1) \qquad \frac{\left(r(X)\right)^{n-1}}{\|X^{n-1}\|_2}$$

is minimized by a bidiagonal matrix whose field of values is a disk. We discuss this matrix, which we call the Crabb matrix, in Section 4. Note that setting $k = n$ in (1.1) implies that the ratio given above is bounded below by $\frac{1}{2}$; this value is attained by the Crabb matrix.

5

3 Subgradients and partial smoothness

We begin with some basic convex analysis on Euclidean spaces — finite-dimensional real inner-product spaces. We consider several specific Euclidean spaces. For integers $n \geq 1$, we consider the space $\mathbf{C}^n$ of complex column n-vectors with inner product $\langle x, y \rangle = \mathrm{Re}(x^*y)$, the space of n-by-n complex matrices $\mathbf{M}^n$ with the inner product $\langle X, Y \rangle = \mathrm{Re}\,\mathrm{trace}(X^*Y)$, and the subspace $\mathbf{H}^n$ of Hermitian matrices (with the same inner product). Unless otherwise stated, $\| \cdot \|$ always refers to the norm induced by the inner product. Dimension always refers to the real vector space notion. A set of vectors $\{x_k : k \in K\}$ is affinely independent in a Euclidean space $\mathbf{E}$ if and only if the vectors $\{(x_k, 1) : k \in K\}$ are linearly independent in the product space $\mathbf{E} \times \mathbf{R}$. The notation ri denotes relative interior.

Given a Euclidean space $\mathbf{E}$, and a norm ν on $\mathbf{E}$, a vector $y \in \mathbf{E}$ is, by definition, a subgradient of ν at a point $x \in \mathbf{E}$ when

$$(3.1) \qquad \langle y, w - x \rangle \leq \nu(w) - \nu(x) \quad \text{for all } w \in \mathbf{E}.$$

We write the set of subgradients as $\partial \nu(x)$.

Given a nonempty convex set $C \subset \mathbf{E}$, the subspace parallel to C is defined by

$$\mathrm{par}\, C \;=\; \mathbf{R}(C - C).$$

By contrast, the linear span of C is

$$\mathrm{span}\, C \;=\; \mathbf{R}C - \mathbf{R}C \;\supset\; \mathrm{par}\, C.$$

The two subspaces are identical if and only if $0 \in C$.

Proposition 3.2 *Consider any norm ν on a Euclidean space $\mathbf{E}$. At any point $x \in \mathbf{E}$, we have the following relationships:*

$$
\begin{aligned}
\partial \nu(x) &= \big\{ y \in \partial \nu(0) : \langle y, x \rangle = \nu(x) \big\} \\
\mathrm{par}\big(\partial \nu(x)\big) &= x^\perp \cap \mathrm{span}\big(\partial \nu(x)\big).
\end{aligned}
$$

Proof The first equation follows from a standard argument, as follows. By the definition of subgradients, we have $y \in \partial \nu(0)$ if and only if

$$(3.3) \qquad \langle y, w \rangle \leq \nu(w) \quad \text{for all } w \in \mathbf{E}.$$

If such a subgradient y also satisfies $\langle y, x \rangle = \nu(x)$, then subtracting implies inequality (3.1). Conversely, inequality (3.1) implies $\langle y, x \rangle = \nu(x)$, by setting w equal to $2x$ and 0 in turn, and then adding implies inequality (3.3).

Turning to the second equation, any element of the left-hand side can be written in the form $u = \mu(y - z)$ for a scalar $\mu \in \mathbf{R}$ and subgradients $y, z \in \partial \nu(x)$. Clearly

u lies in span$(\partial\nu(x))$, and the first equation guarantees $\langle u, x\rangle = 0$, so u also lies in the right-hand side.

For the converse, since $0 \in \partial\nu(0)$, there is nothing to prove when $x = 0$, so we can assume $x \neq 0$, and hence $\nu(x) \neq 0$. Any element of the right-hand side can be written in the form $u = \mu y - \tau z$ for a scalars $\mu, \tau \in \mathbf{R}$ and subgradients $y, z \in \partial\nu(x)$, where $\langle\mu y - \tau z, x\rangle = 0$. Using the first equation, we deduce $(\mu - \tau)\nu(x) = 0$ so $\mu = \tau$, and the result follows. $\qquad\square$

Following [24], we call a norm (or, more generally, an extended-valued convex function) ν on $\mathbf{E}$ *partly smooth* at a point $x \in \mathbf{E}$ relative to a set $\mathcal{M} \subset \mathbf{E}$ containing x if $\mathcal{M}$ is in fact a $\mathcal{C}^{(2)}$-smooth manifold around x whose normal space satisfies $N_{\mathcal{M}}(x) = \operatorname{par}(\partial\nu(x))$, and such that the restriction of ν to $\mathcal{M}$ is $\mathcal{C}^{(2)}$-smooth and the subdifferential mapping $\partial\nu$ restricted to $\mathcal{M}$ is continuous. In particular, any norm is partly smooth at zero relative to $\{0\}$. This definition captures the smooth/nonsmooth geometry of the norm ν around the point x. Specifically, the convex function $y \mapsto \nu(x + y)$ restricted to the tangent space $T_{\mathcal{M}}(x)$ is differentiable at zero, but is *sharp* in nonzero directions $y \notin T_{\mathcal{M}}(x)$: the sum of the directional derivatives in the directions y and $-y$ is strictly positive.

For example, for any point $x \in \mathbf{R}^n$, defined the index set $J = \{i : x_j = 0\}$. Then the l_1-norm $\|\cdot\|_1$ is partly smooth at x relative to the manifold

$$\{z \in \mathbf{R}^n : z_j = 0 \text{ for all } j \in J\}.$$

Then the l_1-norm is partly smooth at any point in $\mathcal{M}$ relative to $\mathcal{M}$.

When a norm is partly smooth relative to a set, we can view that set locally as the set of solutions of a parametrized family of optimization problems. To illustrate, consider the proximal operator corresponding to a proper closed convex function ν on $\mathbf{E}$, namely the map $\operatorname{prox}_\nu \colon \mathbf{E} \to \mathbf{E}$ defined by

$$\operatorname{prox}_\nu(u) \;=\; \operatorname*{argmin}_{x \in \mathbf{E}} \left\{\nu(x) + \frac{1}{2}\|x - u\|^2\right\} \qquad (u \in \mathbf{E}).$$

Proximal operators are central to many first-order algorithms for convex optimization. They are often tractable to compute: in particular, if the function ν is semidefinite representable, as is the case for the numerical radius (by Theorem 1.2), then we can compute the corresponding proximal operator by semidefinite programming.

Among several useful consequences of the idea of partial smoothness, particularly intuitive is its use in sensitivity analysis, by which we mean the behavior of solutions to optimization problems under small perturbations to the data. Under reasonable conditions, the manifold in the definition consists, locally, of all possible optimal solutions to perturbed problems.

As a typical example, we consider the prox operator associated with a convex function ν. (See [25, Cor 4.8], and [13] for more general results along the same lines.)

Note that y is a subgradient of ν at a point x if and only if prox_ν maps $x + y$ to x. Partial smoothness reveals what happens under perturbations.

Theorem 3.4 (Proximal operators and partial smoothness)
Suppose a proper closed convex function ν on a Euclidean space $\mathbf{E}$ is partly smooth at a point $x \in \mathbf{E}$ relative to a set $\mathcal{M} \subset \mathbf{E}$. Consider any subgradient $y \in \mathrm{ri}\big(\partial\nu(x)\big)$. Then the proximal operator prox_ν maps $x + y$ to x, and maps $x + N$, where N is any small neighborhood of y, onto a neighborhood of x in $\mathcal{M}$.

In particular, with the assumptions of this theorem, with positive probability, the proximal operator prox_ν maps vectors chosen randomly from a probability distribution with support $\mathbf{E}$ into the set $\mathcal{M}$. Thus optimization problems (and proximal operators) commonly produce solutions in the manifold $\mathcal{M}$, even though its dimension may be low. The results we prove below about partial smoothness of the numerical radius relative to the set of disk matrices suggest that analogous behavior may be at play in the numerical examples we presented in Section 2.

4　The Crabb matrix

In this work, we prove partial smoothness of the numerical radius around some interesting matrices. Before proceeding to our general development, we illustrate our results on some simple special cases. Here and throughout the paper, we denote the unit circle in the complex plane by $\mathbf{T}$.

The two-by-two case

In $\mathbf{M}^2$, disk matrices coincide with nilpotent matrices, as follows immediately from [10]. In particular, consider the Jordan block

$$J = \begin{bmatrix} 0 & 2 \\ 0 & 0 \end{bmatrix}$$

scaled so that its field of values is the unit disk. As shown in [17, p. 231 and p. 238], at J, the numerical radius r has subdifferential

$$\partial r(J) \; = \; \frac{1}{2}\mathrm{conv}\Big\{ \begin{bmatrix} w & 1 \\ w^2 & w \end{bmatrix} : w \in \mathbf{T} \Big\},$$

and furthermore, integrating over $w \in \mathbf{T}$ shows $\frac{1}{4}J \in \mathrm{ri}\big(\partial r(J)\big)$ (see [17, p. 242]). As we show later, r is partly smooth at J relative to the set of disk matrices $\mathcal{A}$, or equivalently in the case $n = 2$, the set of nilpotent matrices. Hence, by Theorem 3.4, the proximal operator for the numerical radius, prox_r, maps all matrices near $\frac{5}{4}J$ to nilpotent matrices near J. Computational tests, using semidefinite programming, confirm this property.

The three-by-three case

The analogous behavior in $\mathbf{M}^3$ is also simple to check. Now we consider the Jordan block

$$J = \begin{bmatrix} 0 & \sqrt{2} & 0 \\ 0 & 0 & \sqrt{2} \\ 0 & 0 & 0 \end{bmatrix},$$

again scaled so that its field of values is the unit disk. Following the same argument as in [17, p. 231], it is easy to see that at J, the numerical radius r has subdifferential

$$\partial r(J) = \frac{1}{4} \mathrm{conv}\left\{ \begin{bmatrix} w & \sqrt{2} & w^{-1} \\ \sqrt{2}w^2 & 2w & \sqrt{2} \\ w^3 & \sqrt{2}w^2 & w \end{bmatrix} : w \in \mathbf{T} \right\}.$$

Integrating again, we deduce $\frac{1}{4}J \in \mathrm{ri}\big(\partial r(J)\big)$. As we show later, r is partly smooth at J relative to $\mathcal{A}$, so the proximal operator prox_r maps all matrices near $\frac{5}{4}J$ to disk matrices near J. Computational tests again confirm this property.

As we shall see (Theorem 9.10), a necessary condition for $X \in \mathbf{M}^3$ to be a disk matrix is:

X is unitarily similar to a matrix of the form

$$\begin{bmatrix} 0 & a & b \\ 0 & c & d \\ 0 & 0 & 0 \end{bmatrix}$$

for numbers $a, b, c, d \in \mathbf{C}$ satisfying the condition $c(|a|^2 + |d|^2) = -adb^*$.

In fact this condition along with the condition that a and d are both nonzero is also sufficient for $X \in \mathcal{A}$ (by the same result). For X near J, these necessary and sufficient conditions are equivalent, and furthermore easy to check computationally by reduction to Schur form. We show (Theorem 9.13) that the numerical radius is partly smooth at all such matrices, relative to $\mathcal{A}$.

By contrast, we argue at the end of this work (Proposition 9.15) that the numerical radius r is *not* partly smooth relative to the set of disk matrices $\mathcal{A}$ (or indeed any subset) at the disk matrix

$$E_0 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Our later results imply $\frac{1}{2}E_0 \in \mathrm{ri}\big(\partial r(E_0)\big)$, so certainly we have $\mathrm{prox}_r\big(\frac{3}{2}E_0\big) = E_0$. However, the lack of partial smoothness means that there is no guarantee that the proximal operator prox_r maps matrices near $\frac{3}{2}E_0$ to disk matrices.

The general case

Extending the three-by-three case to $\mathbf{M}^n$ for $n > 3$ involves not a Jordan block but the Crabb matrix [8, 19][2]

$$
K \;=\;
\begin{bmatrix}
0 & \sqrt{2} & & & & \\
 & \cdot & 1 & & & \\
 & & \cdot & \cdot & & \\
 & & & \cdot & \cdot & \\
 & & & & \cdot & 1 \\
 & & & & & \cdot & \sqrt{2} \\
 & & & & & & 0
\end{bmatrix}.
$$

This matrix has field of values the unit disk, and again following the reasoning in [17, p. 231], the matrix

$$
H \;=\; \frac{1}{n-1}
\begin{bmatrix}
0 & \frac{1}{\sqrt{2}} & & & & \\
 & \cdot & 1 & & & \\
 & & \cdot & \cdot & & \\
 & & & \cdot & \cdot & \\
 & & & & \cdot & 1 \\
 & & & & & \cdot & \frac{1}{\sqrt{2}} \\
 & & & & & & 0
\end{bmatrix}
$$

is a subgradient of the numerical radius r at K, in fact belonging to $\mathrm{ri}\big(\partial r(K)\big)$. As we show later, r is partly smooth at K relative to the disk matrices $\mathcal{A}$, so the proximal operator prox_r maps all matrices near $K + H$ to disk matrices near K. Again, we can verify this behavior computationally.

As we shall see, a necessary condition for a matrix X to be a disk matrix (by Proposition 5.3) is that the maximal eigenvalues $\lambda_{\max}(w^*X + wX^*)$, as w ranges over any subset of $\mathbf{T}$ of cardinality $2n + 1$, are all equal. This necessary condition is easy to check computationally, and in fact it is also sufficient for X near K (by Proposition 7.1).

---

[2]This matrix was called the "Choi-Crouzeix" matrix in [17, 18]. Many thanks to A. Salemi for providing the reference [8] and pointing out that as a consequence of Crabb's results, the conjectures on [18, p. 239] and [18, p. 240] are true. A third conjecture on [18, p. 242] remains open.

5 The subdifferential of the numerical radius at disk matrices

Clearly both the field of values $W(X)$ and the numerical radius $r(X)$ are invariant under unitary similarity transformations on the matrix X. Since we can write

$$(5.1) \qquad r(X) \;=\; \max \big\{ \mathrm{Re}(w^*u^*Xu) : w \in \mathbf{C}, \; |w| = 1, \; u \in \mathbf{C}^n, \|u\| = 1 \big\},$$

it follows immediately that the function r is a norm.

Our analysis depends heavily on the "Hermitian part" linear map $\Phi \colon \mathbf{M}^n \to \mathbf{H}^n$, defined by

$$\Phi(X) = \frac{1}{2}(X + X^*) \qquad (X \in \mathbf{M}^n).$$

Notice that this map is the adjoint of the natural embedding: $\Phi^* \colon \mathbf{H}^n \to \mathbf{M}^n$ is given by

$$(5.2) \qquad\qquad\qquad \Phi^*(Z) = Z \qquad (Z \in \mathbf{H}^n),$$

since

$$\begin{aligned}
\langle X, \Phi^*(Z) \rangle \;&=\; \langle \Phi(X), Z \rangle \;=\; \Big\langle \frac{1}{2}(X + X^*), Z \Big\rangle \;=\; \frac{1}{2}\big(\langle X, Z \rangle + \langle X^*, Z \rangle\big) \\
&=\; \frac{1}{2}\big(\mathrm{Re}\,\mathrm{tr}(X^*Z) + \mathrm{Re}\,\mathrm{tr}(XZ)\big) \;=\; \frac{1}{2}\big(\mathrm{Re}\,\mathrm{tr}(X^*Z) + \mathrm{Re}\,\mathrm{tr}(ZX^*)\big) \\
&=\; \frac{1}{2}\big(\mathrm{Re}\,\mathrm{tr}(X^*Z) + \mathrm{Re}\,\mathrm{tr}(X^*Z)\big) \;=\; \langle X, Z \rangle.
\end{aligned}$$

We will be interested in the largest eigenvalue $\lambda_{\max}$ of matrices of the form $\Phi(w^*X)$ for scalars w on the unit circle $\mathbf{T} \subset \mathbf{C}$. We use the following result (dating back to [21]) extensively: since its proof is simple, we include it.

Proposition 5.3 (Disk matrix characterization) *For any matrix $X \in \mathbf{M}^n$,*

$$\max \big\{ \langle w, z \rangle : z \in W(X) \big\} \;=\; \lambda_{\max}\big(\Phi(w^*X)\big) \;\; \textit{for all } w \in \mathbf{T}.$$

Hence X has numerical radius

$$r(X) = \max_{w \in \mathbf{T}} \lambda_{\max}\big(\Phi(w^*X)\big),$$

and has field of values a disk of radius $\rho \geq 0$ if and only if $\lambda_{\max}\big(\Phi(w^*X)\big) = \rho$ for all $w \in \mathbf{T}$.

Proof Consider the first equation. The field of values is nonempty and compact, so the maximum on the left-hand side is indeed attained. Furthermore we have

$$
\begin{aligned}
& \max \left\{ \langle w, z \rangle : z \in W(X) \right\} \\
&= \max \left\{ \operatorname{Re}(w^* u^* X u) : u \in \mathbf{C}^n, \|u\| = 1 \right\} \\
&= \max \left\{ \frac{1}{2}(w^* u^* X u + u^* X^* u w) : u \in \mathbf{C}^n, \|u\| = 1 \right\} \\
&= \max \left\{ u^* \left(\frac{1}{2}(w^* X + w X^*) \right) u : u \in \mathbf{C}^n, \|u\| = 1 \right\} \\
&= \lambda_{\max} \left(\Phi(w^* X) \right),
\end{aligned}
$$

so the first equation follows. The second equation is an immediate consequence. One direction in the final equivalence is also immediate, and the converse follows since the field of values $W(X)$ is convex by the Toeplitz-Hausdorff Theorem. $\square$

Lemma 5.4 *Given a matrix $\hat{X} \in \mathbf{M}^n$ and a scalar $w \in \mathbf{T}$, suppose that the matrix $\Phi(w^* \hat{X})$ has simple largest eigenvalue, with a corresponding unit eigenvector g. Then the function $X \mapsto \lambda_{\max} \left(\Phi(w^* X) \right)$ is analytic on $\mathbf{M}^n$ around $\hat{X}$, and its gradient there is wgg^*. Furthermore, the following condition holds:*

$$
\langle wgg^*, \hat{X} \rangle = \lambda_{\max} \left(\Phi(w^* \hat{X}) \right).
$$

Proof Since the largest eigenvalue of the matrix $\Phi(w^* \hat{X})$ is simple, it is well known [35] that the function $\lambda_{\max} \colon \mathbf{H}^n \to \mathbf{R}$ is locally analytic there, with gradient gg^*. Therefore, for small matrices $Z \in \mathbf{H}^n$ we have

$$
\begin{aligned}
& \lambda_{\max} \left(\Phi \left(w^* (\hat{X} + Z) \right) \right) - \lambda_{\max} \left(\Phi(w^* \hat{X}) \right) \\
&= \lambda_{\max} \left(\Phi(w^* \hat{X}) + \Phi(w^* Z) \right) - \lambda_{\max} \left(\Phi(w^* \hat{X}) \right) \\
&= \langle gg^*, \Phi(w^* Z) \rangle + o(\|Z\|).
\end{aligned}
$$

Using the adjoint property (5.2), notice

$$
\langle gg^*, \Phi(w^* Z) \rangle = \langle \Phi^*(gg^*), w^* Z \rangle = \operatorname{Re} \operatorname{tr}\left(gg^* w^* Z \right) = \langle wgg^*, Z \rangle.
$$

The expression for the gradient follows. We then note

$$
\begin{aligned}
\langle wgg^*, \hat{X} \rangle &= \operatorname{Re} \operatorname{tr}\left(w^* gg^* \hat{X} \right) = \operatorname{Re}\left(g^* w^* \hat{X} g \right) \\
&= g^* \left(\Phi(w^* \hat{X}) \right) g = g^* \lambda_{\max} \left(\Phi(w^* \hat{X}) \right) g = \lambda_{\max} \left(\Phi(w^* \hat{X}) \right),
\end{aligned}
$$

using the definition of g. $\square$

We next define a set-valued mapping $\Delta \colon \mathbf{M}^n \rightrightarrows \mathbf{M}^n$ by

$$(5.5) \qquad \Delta(X) \;=\; \big\{ wgg^* : w \in \mathbf{T}, \; \|g\| = 1, \; \Phi(w^*X)g = \lambda_{\max}\big(\Phi(w^*X)\big)g \big\},$$

and a set of matrices

$$(5.6) \qquad \Omega \;=\; \big\{ X \in \mathbf{M}^n : \lambda_{\max}\big(\Phi(w^*X)\big) \text{ is simple for all } w \in \mathbf{T} \big\}$$

Proposition 5.7 *On the open set Ω, the set-valued mappings $X \mapsto \Delta(X)$ and $X \mapsto \mathrm{conv}\big(\Delta(x)\big)$ are both continuous.*

Proof Suppose the set Ω is not open, so there exists a matrix $X \in \Omega$, and a sequence $X_r \to X$ with $X_r \notin \Omega$ for all $r = 1, 2, 3, \ldots$. By definition, for each r there exists a scalar $w_r \in \mathbf{T}$ such that the second largest eigenvalue $\lambda_2\big(\Phi(w_r^*X_r)\big)$ equals the largest $\lambda_{\max}\big(\Phi(w_r^*X_r)\big)$. After taking a subsequence, we can suppose w_r converges to some scalar $w \in T$, but then, since the map Φ and both the eigenvalues $\lambda_{\max}$ and λ_2 are continuous, we deduce the contradiction

$$\lambda_2\big(\Phi(w^*X)\big) = \lambda_{\max}\big(\Phi(w^*X)\big).$$

Hence Ω is indeed open.

To prove continuity, we use Lemma 5.4, which shows that, for any scalar $w \in \mathbf{T}$, the function $F \colon \mathbf{M}^n \to \mathbf{R}$ defined by $F(X) = \lambda_{\max}\big(\Phi(w^*X)\big)$ is analytic throughout the set Ω. Furthermore, the gradient of F at any matrix $X \in \Omega$ is given by $\nabla F(X) = wgg^*$ for any unit eigenvector g corresponding to the eigenvalue $\lambda_{\max}\big(\Phi(w^*X)\big)$. Now consider a sequence $X_r \to X$ in Ω. We prove $\Delta(X_r) \to \Delta(X)$ in two steps.

First, we observe that any limit point Y of a sequence $Y_r \in \Delta(X_r)$ must lie in $\Delta(X)$. To see this, after taking a subsequence we can suppose $Y_r \to Y$, and $Y_r = w_r g_r g_r^*$ for some scalar $w_r \in \mathbf{T}$ and unit vector g_r satisfying

$$\Phi(w_r^*X_r)g_r \;=\; \lambda_{\max}\big(\Phi(w_r^*X_r)\big)g_r \quad \text{for all } r.$$

Taking a further subsequence, we can suppose w_r approaches some scalar $w \in \mathbf{T}$ and g_r approaches some unit vector g. By continuity we deduce

$$Y = wgg^* \quad \text{and} \quad \Phi(w^*X)g \;=\; \lambda_{\max}\big(\Phi(w^*X)\big)g,$$

so we have shown $Y \in \Delta(X)$, as desired.

Secondly, given any matrix $Y \in \Delta(X)$, we prove the existence of a sequence $Y_r \in \Delta(X_r)$ converging to Y. To do so, first write $Y = wgg^*$ for some unit eigenvector g corresponding to the eigenvalue $\lambda_{\max}\big(\Phi(w^*X)\big)$, so $Y = \nabla F(X)$. If we define $Y_r = \nabla F(X_r)$ for each r, then analyticity implies $Y_r \to Y$, and applying Lemma 5.4 at each X_r shows $Y_r \in \Delta(X_r)$, as desired.

Continuity of the map Φ and the eigenvalue $\lambda_{\max}$ shows that, for any matrix $X \in \Omega$, the image $\Delta(X)$ is always closed, and clearly every element has norm one,

so the mapping Δ is also locally (in fact globally) bounded. Continuity of the mapping $X \mapsto \text{conv}\big(\Delta(x)\big)$ now follows in a standard fashion [33, Prop. 4.30]. □

Theorem 5.8 (Disk matrix subgradients) *At any matrix X in the set $\mathcal{A}$ of matrices whose fields of values are disks centered at zero, the subdifferential of the numerical radius r satisfies*

$$\partial r(X) = \text{conv}\big(\Delta(X)\big)$$
$$\text{par}\big(\partial r(X)\big) = X^{\perp} \cap \text{span}\big(\Delta(X)\big),$$

where the set-valued mapping Δ is defined by equation (5.5). On the intersection of $\mathcal{A}$ with the set Ω defined by equation (5.6), the numerical radius is analytic and its subdifferential is continuous.

Proof We can write equation (5.1) as

$$r(X) = \max_{w \in \mathbf{T}} \max_{\|g\|=1} \langle wgg^*, X \rangle,$$

for all matrices $X \in \mathbf{M}^n$, so

$$\partial r(X) = \text{conv}\big\{ wgg^* : w \in \mathbf{T},\ \|g\| = 1,\ \langle wgg^*, X \rangle = r(X) \big\}$$

(see [33, 10.31]). For any fixed scalar $w \in \mathbf{T}$, the inner maximum in the expression for the numerical radius equals the eigenvalue $\lambda_{\max}\big(\Phi(w^*X)\big)$, and the attaining set consists of all corresponding unit eigenvectors g. For any matrix $X \in \mathcal{A}$, the outer maximum is attained on the whole set $\mathbf{T}$, and hence $\partial r(X) = \text{conv}\big(\Delta(X)\big)$. The second equation follows from Proposition 3.2. Continuity follows by Proposition 5.7, and since for all $X \in \mathcal{A}$ we have $r(X) = \lambda_{\max}\big(\Phi(X)\big)$, analyticity follows from Lemma 5.4. □

6   Matrices with field of values the unit disk

Let $\mathcal{B}$ denote the set of matrices with field of values the *unit* disk; hence $\mathcal{A} = \mathbf{R}_+\mathcal{B}$. In this section we describe conditions under which $\mathcal{B}$ is locally a manifold. We focus on the subset $\mathcal{B} \cap \Omega$, consisting of those matrices $X \in \mathcal{B}$ for which the largest eigenvalue of the matrix $\Phi(w^*X)$ equals one and is simple for all scalars $w \in \mathbf{T}$. We begin with a tool for characterizing when a matrix is in $\mathcal{B}$, or equivalently, in $\mathcal{B} \cap \Omega$, given that it is close to a given matrix in $\mathcal{B} \cap \Omega$.

Proposition 6.1 *Consider any matrix $\hat{X}$ in the set $\mathcal{B} \cap \Omega \subset \mathbf{M}^n$. Then for any matrix X near $\hat{X}$, and any set $\hat{\mathbf{T}} \subset \mathbf{T}$ of cardinality at least $2n + 1$, the following properties are equivalent.*

(i) $X \in \mathcal{B}$.

(ii) $\lambda_{\max}\big(\Phi(w^*X)\big) = 1$ *for all* $w \in \hat{\mathbf{T}}$.

(iii) $\Phi(w^*X) - I$ *is singular for all* $w \in \hat{\mathbf{T}}$.

(iv) $\Phi(w^*X) - I$ *is singular for all* $w \in \mathbf{T}$.

(v) $\lambda_{\max}\big(\Phi(w^*X)\big) = 1$ *for all* $w \in \mathbf{T}$.

Proof The implications (v) $\Rightarrow$ (i) $\Rightarrow$ (ii) follow from Proposition 5.3. The impli-
cation (ii) $\Rightarrow$ (iii) is immediate.

Notice the identity

$$w^n \det\big(\Phi(w^*X) - I\big) \;=\; \det\left(\tfrac{1}{2}(X + w^2 X^*) - wI\right) \quad \text{for all } w \in \mathbf{T}.$$

The right-hand side is a polynomial in w of degree at most $2n$, so it vanishes ev-
erywhere on the set $\hat{\mathbf{T}}$ if and only if it vanishes everywhere on $\mathbf{T}$. It follows that
properties (iii) and (iv) are equivalent.

Suppose the implication (iv) $\Rightarrow$ (v) fails. Then there exists a sequence of matrices
X_r converging to $\hat{X}$, and scalars $w_r \in \mathbf{T}$, such that 1 is an eigenvalue of the matrix
$\Phi(w_r^*X_r)$ and yet is not its largest eigenvalue. Consequently, its second largest
eigenvalue λ_2 satisfies $\lambda_2\big(\Phi(w_r^*X_r)\big) \geq 1$ for all r. After taking a subsequence, we
can suppose $w_r \to w \in \mathbf{T}$, and we deduce by continuity $\lambda_2\big(\Phi(w^*\hat{X})\big) \geq 1$. But this
contradicts our assumption that the matrix $\Phi(w^*\hat{X})$ has a simple largest eigenvalue
equal to 1. $\qquad\square$

We now define functions $f \colon \mathbf{T} \to \mathbf{C}^n$ and $F \colon \mathbf{T} \to \mathbf{M}^n$ by

$$f(w) = \begin{bmatrix} w \\ w^2 \\ w^3 \\ \vdots \\ w^n \end{bmatrix}$$

and

$$F(w) \;=\; wf(w)f(w)^* \;=\; \begin{bmatrix} w & 1 & w^{-1} & w^{-2} & \cdots & w^{2-n} \\ w^2 & w & 1 & w^{-1} & \cdots & w^{3-n} \\ w^3 & w^2 & w & 1 & \cdots & w^{4-n} \\ w^4 & w^3 & w^2 & w & \cdots & w^{5-n} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \ddots \\ w^n & w^{n-1} & w^{n-2} & w^{n-3} & \cdots & w \end{bmatrix}$$

15

Proposition 6.2 *For any set $\tilde{\mathbf{T}} \subset \mathbf{T}$ of cardinality at most $2n+1$, the set of vectors $\{f(w) : w \in \tilde{\mathbf{T}}\}$ is affinely independent in the Euclidean space $\mathbf{C}^n$, and the set of matrices $\{F(w) : w \in \tilde{\mathbf{T}}\}$ is linearly independent in the Euclidean space $\mathbf{M}^n$.*

Proof We begin with the first claim. Suppose that it fails for some set $\tilde{\mathbf{T}}$ of distinct scalars $w_1, w_2, w_3, \ldots, w_m$ where $m \leq 2n + 1$. Without loss of generality, we can assume $m = 2n + 1$: otherwise, simply augment the set $\tilde{\mathbf{T}}$.

Then the linear map $\Psi \colon \mathbf{R}^{2n+1} \to \mathbf{R} \times \mathbf{C}^n$ defined by

$$\Psi(y) \;=\; \left(\sum_j y_j, \Big(\sum_{j=1}^{2n+1} y_j w_j^k \Big)_{k=1}^{n} \right) \quad (y \in \mathbf{R}^{2n+1})$$

is not one-to-one. Since $\dim \mathbf{R}^{2n+1} = 2n+1 = \dim(\mathbf{R} \times \mathbf{C}^n)$, the map Ψ is not onto either, and hence its adjoint $\Psi^* \colon \mathbf{R} \times \mathbf{C}^n \to \mathbf{R}^{2n+1}$ is not one-to-one. In other words, there exists a scalar $\alpha \in \mathbf{R}$ and $x \in \mathbf{C}^n$, not both zero, such that $\Psi^*(\alpha, x) = 0$. By definition, for all vectors $y \in \mathbf{R}^{2n+1}$ we have

$$\big\langle \Psi^*(\alpha, x), y \big\rangle \;=\; \big\langle (\alpha, x), \Psi(y) \big\rangle$$
$$= \; \alpha \sum_{j=1}^{2n+1} y_j + \mathrm{Re}\Big(\sum_{k=1}^{n} x_k^* \sum_{j=1}^{2n+1} y_j w_j^k \Big) \;=\; \sum_{j=1}^{2n+1} y_j \Big(\alpha + \mathrm{Re} \sum_{k=1}^{n} x_k^* w_j^k \Big)$$

so for $j = 1, 2, 3, \ldots, 2n + 1$ we have $0 = \big(\Psi^*(\alpha, x) \big)_j = \alpha + \mathrm{Re} \sum_k x_k^* w_j^k$. However, the nonzero real trigonometric polynomial $\theta \mapsto \alpha + \mathrm{Re} \sum_k x_k^* e^{ik\theta}$ can have at most $2n$ zeroes [32, p. 150], which is a contradiction.

This completes the proof of the first claim. The second claim follows immediately, since, for all $w \in \mathbf{T}$, the first column of the matrix $F(w)$ is the vector $f(w)$, and the $(1, 2)$ entry equals 1. $\qquad \square$

Corollary 6.3 *In the Euclidean space $\mathbf{C}^{n+1}$, the convex set*

$$\mathrm{conv}\Big\{ \begin{bmatrix} 1 \\ f(w) \end{bmatrix} : w \in \mathbf{T} \Big\}$$

has dimension $2n$.

Proof The given set, which we denote by S, contains the set

$$\Big\{ \begin{bmatrix} 1 \\ f(w) \end{bmatrix} : w \in \hat{\mathbf{T}} \Big\}$$

for any set $\hat{\mathbf{T}} \subset \mathbf{T}$. By choosing $\hat{\mathbf{T}}$ of cardinality exactly $2n+1$, we see by Proposition 6.2 that S contains $2n+1$ linearly independent vectors, and hence has dimension at

least $2n$. On the other hand, $y_1 = 1$ for all vectors $y \in S$, so S can have dimension no more than $2n$. $\quad\square$

To prove that the set $\mathcal{B}$ of matrices with field of values the unit disk is a manifold locally, we focus on matrices $X \in \mathcal{B} \cap \Omega$, and make one further assumption: that the matrix $\Phi(w^*X)$ parametrized by $w \in \mathbf{T}$ has an (unnormalized) eigenvector corresponding to the simple eigenvalue $\lambda_{\max}\big(\Phi(w^*X)\big) = 1$ whose components are linearly independent polynomials in w of degree less than n. At first sight, this assumption seems artificial, but in the sections that follow we provide motivation and examples.

Theorem 6.4 *Consider any matrix $X \in \mathbf{M}^n$. Suppose there exists an invertible matrix $G \in \mathbf{M}^n$ such that the largest eigenvalue of the matrix $\Phi(w^*X)$ is simple and equal to one for all scalars $w \in \mathbf{T}$, with a corresponding eigenvector $Gf(w)$. Then $X \in \mathcal{B}$ and the set $\mathcal{B}$ of matrices with field of values the unit disk is an analytic manifold around X of codimension $2n + 1$. Its tangent and normal spaces at X are*

$$
\begin{aligned}
T_{\mathcal{B}}(X) &= \Big\{ Y \in \mathbf{M}^n : \langle GF(w)G^*, Y \rangle = 0 \text{ for all } w \in \mathbf{T} \Big\} \\
N_{\mathcal{B}}(X) &= G\Big(\operatorname{span}\{F(w) : w \in \mathbf{T}\} \Big) G^*.
\end{aligned}
$$

Proof That $X \in \mathcal{B}$ is immediate from Proposition 6.1. For each scalar $w \in \mathbf{T}$, denote the normalized eigenvector $\|Gf(w)\|^{-1}Gf(w)$ by $g(w)$. Choose any any set $\hat{\mathbf{T}} \subset \mathbf{T}$ of cardinality exactly $2n + 1$. By Proposition 6.1, the set $\mathcal{B}$ is described locally by the analytic equations in its property (ii). We see from Lemma 5.4 that the gradients of the left-hand sides of these equations are linearly independent, since for any $y \in \mathbf{R}^{2n+1}$, the relationship

$$
\sum_j y_j w_j g(w_j) g(w_j)^* = 0
$$

implies

$$
\begin{aligned}
0 &= \sum_j y_j w_j \|Gf(w_j)\|^{-2} Gf(w_j) f(w_j)^* G^* \\
&= G\Big(\sum_j y_j \|Gf(w_j)\|^{-2} \big(w_j f(w_j) f(w_j)^*\big) \Big) G^* \\
&= G\Big(\sum_j y_j \|Gf(w_j)\|^{-2} F(w_j) \Big) G^*,
\end{aligned}
$$

and hence

$$
\sum_j \big(y_j \|Gf(w_j)\|^{-2}\big) F(w_j) = 0,
$$

17

and then Proposition 6.2 shows $y_j \|Gf(w_j)\|^{-2} = 0$ for all j and so $y = 0$. The manifold property follows and

$$T_{\mathcal{B}}(X) = \left\{ Y \in \mathbf{M}^n : \langle wg(w)g(w)^*, Y \rangle = 0 \text{ for all } w \in \hat{\mathbf{T}} \right\}.$$

Since $\hat{\mathbf{T}}$ was arbitrary, the formula for the tangent space follows, and that for the normal space is an immediate consequence. $\qquad\square$

7　Disk matrices

Our next step is to lift the main result in the previous section from the set $\mathcal{B}$ to the set of disk matrices $\mathcal{A}$. We first note a computational technique for recognizing when a matrix is in $\mathcal{A}$, given that it is close to a given matrix in $\mathcal{A} \cap \Omega$.

Proposition 7.1 (Recognizing disk matrices) *Consider any fixed matrix $\hat{X}$ in the set $\mathcal{A} \cap \Omega \subset \mathbf{M}^n$ and any set $\hat{\mathbf{T}} \subset \mathbf{T}$ of cardinality at least $2n + 1$. A necessary condition for a matrix $X \in \mathbf{M}^n$ to be a disk matrix is that the eigenvalues $\left\{ \lambda_{\max}(\Phi(w^*X)) : w \in \hat{\mathbf{T}} \right\}$ should all be equal. For X near $\hat{X}$, this condition is also sufficient.*

Proof　The case $n = 1$ is an easy exercise, relying on the fact that a scalar $x \in \mathbf{C}$ is zero if and only if $\mathrm{Re}(w^*x)$ takes identical values for at least three distinct scalars $w \in \mathbf{T}$. We can therefore assume $n > 1$.

　The necessary condition is clear from Proposition 5.3. Suppose the converse fails, so there exists a sequence of matrices (X_s) outside $\mathcal{A}$ converging to $\hat{X}$ and a real sequence (α_s) such that

$$\lambda_{\max}(\Phi(w^*X_s)) = \alpha_s \text{ for all } w \in \hat{\mathbf{T}}, \text{ and all } s.$$

Clearly we have $\alpha_s \to r(\hat{X}) > 0$. In particular, we can assume $\alpha_s > 0$ for all s. But now we have

$$\frac{1}{\alpha_s} X_s \to \frac{1}{r(\hat{X})} \hat{X} \in \mathcal{B} \cap \Omega,$$

so we can apply Proposition 6.1 to deduce the contradiction $\frac{1}{\alpha_s} X_s \in \mathcal{B}$ for all large s. $\qquad\square$

　We next establish a simple tool.

Lemma 7.2 *Consider any matrix $\hat{X} \in \mathcal{B}$. For all real $\delta > 0$, the set $\mathcal{A}$ of disk matrices coincides with the set*

$$\mathcal{A}_\delta = \left\{ \rho X : X \in \mathcal{B}, \ \rho \in \mathbf{R}, \ \|X - \hat{X}\| < \delta, \ |\rho - 1| < \delta \right\}$$

around $\hat{X}$.

18

Proof   Since the field of values satisfies $W(\rho X) = \rho W(X)$ for all real $\rho > 0$, we clearly always have $\mathcal{A}_\delta \subset \mathcal{A}$. Conversely, consider any sequence of matrices $X_k \to \hat{X}$ in $\mathcal{A}$. The numerical radius r is continuous, so $r(X_k) \to r(\hat{X}) = 1$. Now note $\frac{1}{r(X_k)} X_k \in \mathcal{B}$, so $X_k \in \mathcal{A}_\delta$ eventually. The result follows. $\square$

Proposition 7.3 *Consider a Euclidean space* $\mathbf{E}$, *and a set* $\mathcal{S} \subset \mathbf{E}$ *that is a k-dimensional analytic manifold around a point* $\bar{s} \in \mathcal{S}$. *Suppose that* $\bar{s}$ *does not lie in the tangent space* $T_\mathcal{S}(\bar{s})$. *Then for all small* $\delta > 0$, *the set*

$$\left\{ \rho s : s \in \mathcal{S}, \ \rho \in \mathbf{R}, \ \|s - \bar{s}\| < \delta, \ |\rho - 1| < \delta \right\}$$

is a (k+1)-dimensional manifold around $\bar{s}$, *and its tangent space there is* $T_\mathcal{S}(\bar{s}) + \mathbf{R}\bar{s}$.

Proof   Denote the open unit ball in the space $\mathbf{E}$ by B. By choosing suitable local coordinates for the manifold $\mathcal{S}$ we can suppose there is an analytic map $P \colon B \to \mathbf{E}$ such that $P(0) = \bar{s}$ and with derivative DP one-to-one throughout B and such that the image $P(B)$ is an open neighborhood of $\bar{s}$ in $\mathcal{S}$. In that case the tangent space $T_\mathcal{S}(\bar{s})$ is just the range of $DP(0)$.

Now define an analytic map $Q \colon B \times \mathbf{R} \to \mathbf{E}$ by $Q(y, \rho) = \rho P(y)$. Notice

$$0 = DQ(0,1)(y, \rho) = DP(0)y + \rho \bar{s}$$

implies $(y, \rho) = (0, 0)$, since by assumption the vector $\bar{s}$ does not lie in the range of $DP(0)$ and the map $DP(0)$ is one-to-one. Therefore the derivative map DQ is also one-to-one at the vector $(0, 1)$, so it is one-to-one throughout the set $\delta B \times (1 - \delta, 1 + \delta)$ for all small $\delta > 0$. The image of this set under the map Q is the given set, so the result now follows. $\square$

We deduce the main result of this section.

Theorem 7.4 (Partial smoothness at disk matrices) *Consider any matrix* X *in* $\mathbf{M}^n$. *Suppose there exists an invertible matrix* $G \in \mathbf{M}^n$ *such that the largest eigenvalue of the matrix* $\Phi(w^* X)$ *is simple and independent of the scalar* $w \in \mathbf{T}$, *with a corresponding eigenvector* $Gf(w)$. *Then around* X, *the set* $\mathcal{A}$ *of disk matrices is an analytic manifold of codimension* $2n$. *Its tangent and normal spaces at* X *are*

$$T_\mathcal{A}(X) = \mathbf{R}X + \left\{ Y \in \mathbf{M}^n : \langle GF(w)G^*, Y \rangle = 0 \text{ for all } w \in \mathbf{T} \right\}$$

$$N_\mathcal{A}(X) = X^\perp \cap G\left(\text{span}\{F(w) : w \in \mathbf{T}\} \right) G^*.$$

On $\mathcal{A}$, *around* X, *the numerical radius is analytic and its subdifferential is continuous and satisfies*

$$\partial r(X) = G\left(\text{conv}\{\|Gf(w)\|^{-2} F(w) : w \in \mathbf{T}\} \right) G^*$$

$$\text{par}(\partial r(X)) = N_\mathcal{A}(X).$$

Hence the numerical radius is partly smooth at the matrix X *relative to* $\mathcal{A}$.

Proof The case $n = 1$ is an easy exercise, so we assume $n > 1$. Suppose first that the largest eigenvalue of $\Phi(w^*X)$ is 1 for all w: in other words, the matrix X has numerical radius 1. Using Theorem 6.4 and the notation of its proof, to apply Proposition 7.3 we just need to check $X \notin \mathbf{T}_\mathcal{B}(X)$. This follows from Lemma 5.4, which shows in fact $\text{Re}\left(w^*g(w)^*Xg(w)\right) = 1$ for all $w \in \mathbf{T}$. The properties of the subdifferential of the numerical radius follow from Theorem 5.8. The result follows.

To prove the general case, define the positive scalar $\alpha = \frac{1}{r(X)}$, and apply the previous argument to the matrix αX. Since the set $\mathcal{A}$ is closed under scalar multiplication, it is easy to check $N_\mathcal{A}(\alpha X) = N_\mathcal{A}(X)$. Furthermore, since the numerical radius is positively homogeneous, it is also easy to check $\partial r(\alpha X) = \partial r(X)$. This completes the proof. □

8   The Crouzeix parametrization

According to Proposition 5.3, a matrix $X \in \mathbf{M}^n$ belongs to the set $\mathcal{B}$ of matrices with field of values the unit disk if and only if the Hermitian matrix $\Phi(w^*X)$ has largest eigenvalue equal to one for all scalars $w \in \mathbf{T}$. To guarantee that $\mathcal{B}$ is a manifold around X, Theorem 6.4 first made a modest refinement of the basic assumption: that the largest eigenvalue is also simple for all $w \in \mathbf{T}$. Secondly, however, we made a more opaque assumption: that we can associate with this eigenvalue a corresponding eigenvector whose components are linearly independent polynomials in w of degree less than n. We now explore this assumption, using a result from [10]. A somewhat related characterization appears in [30, Thm 3.1],

Theorem 8.1 (Crouzeix's disk matrix parametrization) *A matrix in $\mathbf{M}^n$ has field of values the unit disk if and only if it has the form $X = 2(\sin B)U(\cos B)$ for some matrix $B \in \mathbf{H}^n$ satisfying $\frac{\pi}{2}I \succcurlyeq B \succcurlyeq 0$ and some unitary $U \in \mathbf{M}^n$ such that the matrix $U(\cos B) - w(\sin B)$ is singular for all scalars $w \in \mathbf{T}$.*

In this theorem statement, we can write

$$\sin B = V^*(\text{Diag}\, s)V \quad \text{and} \quad \cos B = V^*(\text{Diag}\, c)V,$$

for some unitary $V \in \mathbf{M}^n$ and nonnegative vectors $s, c \in \mathbf{R}^n$ satisfying $s_j^2 + c_j^2 = 1$ for $j = 1, 2, 3, \ldots, n$. With this notation, we have

$$VXV^* = 2(\text{Diag}\, s)(VUV^*)(\text{Diag}\, c)$$

and the matrix $(VUV^*)(\text{Diag}\, c) - w(\text{Diag}\, s))$ is singular for all $w \in \mathbf{T}$. Hence we lose no essential generality in assuming, when convenient, that $V = I$, or in other words that B is diagonal.

Proposition 8.2 *For matrices X as described in Theorem 8.1, the eigenspace corresponding to the largest eigenvalue of the matrix $\Phi(w^*X)$ is just the nullspace of the matrix $U(\cos B) - w(\sin B)$.*

Proof By Theorem 8.1 and Proposition 5.3, we know $\lambda_{\max}(\Phi(w^*X)) = 1$. Consider any vector v in the nullspace of the matrix $U(\cos B) - w(\sin B)$. We deduce

$$w^*(\sin B)U(\cos B)v = (\sin B)^2 v$$
$$w(\cos B)U^*(\sin B)v = (\cos B)^2 v$$

Adding shows $\Phi(w^*X)v = v$, as required.

Conversely, if $\Phi(w^*X)v = v$, then we have

$$\begin{aligned}
\left\| (U(\cos B)v - w(\sin B)v \right\|^2 &= \|U(\cos B)v\|^2 + \|w(\sin B)v\|^2 \\
&\quad - (U(\cos B)v)^*(w(\sin B)v) - (w(\sin B)v)^*(U(\cos B)v) \\
&= \|(\cos B)v\|^2 + \|(\sin B)v\|^2 - v^*\Phi(w^*X)v = 0,
\end{aligned}$$

as required. $\square$

Any matrix with field of values the unit disk can be written in the form described in Theorem 8.1. In that result, the matrix $U(\cos B) - w(\sin B)$ must be singular for all scalars $w \in \mathbf{T}$, or in matrix-analytic terminology, the matrix pencil $U(\cos B) - w(\sin B)$ is singular. Theorem 6.4 needs more: in particular, the largest eigenvalue of the matrix $\Phi(w^*X)$ should be simple, or equivalently (by Proposition 8.2) the nullspace of $U(\cos B) - w(\sin B)$ (in the complex vector space $\mathbf{C}^n$) should be one-dimensional for all $w \in \mathbf{T}$. Assuming this, the nullspace is in general spanned by some vector having all components polynomial in w, as can be seen quickly by reducing the matrix polynomial $U(\cos B) - w(\sin B)$ to Smith normal formal [14]. However, Theorem 6.4 needs more precise knowledge of this eigenvector of $\Phi(w^*X)$. Rather than delving into the theory of singular pencil canonical forms, we simply assume that the n components of this vector can be taken to be linearly independent complex polynomials of degree less than n, or equivalently, that the vector can be written in the form $Gf(w)$, where the matrix $G \in \mathbf{M}^n$ is invertible. This assumption will suffice for the examples we consider.

Theorem 8.3 (Partial smoothness and the Crouzeix parametrization)
Suppose for some matrix $B \in \mathbf{H}^n$ satisfying $\frac{\pi}{2}I \succcurlyeq B \succcurlyeq 0$, and some unitary $U \in \mathbf{M}^n$, there exists an invertible $G \in \mathbf{M}^n$ such that

$$\mathrm{Null}(U(\cos B) - w(\sin B)) = \mathbf{C}Gf(w) \quad \text{for all scalars } w \in \mathbf{T}.$$

Then around the matrix $X = 2(\sin B)U(\cos B)$ the set $\mathcal{A}$ of disk matrices is an analytic manifold of codimension $2n$. Its tangent and normal spaces at X are

$$T_{\mathcal{A}}(X) = \mathbf{R}X + \left\{ Y \in \mathbf{M}^n : \langle GF(w)G^*, Y \rangle = 0 \text{ for all } w \in \mathbf{T} \right\}$$
$$N_{\mathcal{A}}(X) = X^\perp \cap G\left(\mathrm{span}\{F(w) : w \in \mathbf{T}\} \right)G^*.$$

On $\mathcal{A}$, around X, the numerical radius is analytic and its subdifferential is continuous and satisfies

$$\partial r(X) = G\Big(\operatorname{conv}\{\|Gf(w)\|^{-2}F(w) : w \in \mathbf{T}\}\Big)G^*$$
$$\operatorname{par}\big(\partial r(X)\big) = N_{\mathcal{A}}(X).$$

In particular, the numerical radius is partly smooth at the matrix X relative to the manifold $\mathcal{A}$.

Proof This follows directly from Proposition 8.2 and Theorem 7.4. $\square$

We consider two special cases: two-by-two matrices and superdiagonal matrices.

8.1   The two-by-two case

As shown in [10], a two-by-two matrix is a disk matrix if and only if it is unitarily similar to a multiple of the two-by-two Jordan block. In Theorem 8.3, we could take

$$U = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} \frac{\pi}{2} & 0 \\ 0 & 0 \end{bmatrix},$$

so

$$X = 2(\sin B)U(\cos B) = 2\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 2 \\ 0 & 0 \end{bmatrix}$$

and

$$U(\cos B) - w(\sin B) = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} - w\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} -w & 1 \\ 0 & 0 \end{bmatrix}.$$

Thus $G = I$ satisfies the assumptions of the theorem, so the numerical radius is partly smooth at X relative to $\mathcal{A}$.

8.2   The superdiagonal case

More generally, writing matrices in $\mathbf{M}^n$ (for $n > 2$) by partitioning both rows and columns into blocks of size 1, $n - 2$, and 1, consider the case

$$U = \begin{bmatrix} 0 & e_1^* & 0 \\ 0 & J & e_{n-2} \\ 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} \frac{\pi}{2} & 0 & 0 \\ 0 & \operatorname{Diag}\theta & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

where J denotes the $(n-2)$-by-$(n-2)$ Jordan block, $e_1, e_{n-2} \in \mathbf{R}^{n-2}$ are the canonical unit vectors,

$$
J = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & 1 \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}, \quad e_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}, \quad e_{n-2} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}
$$

and the vector $\theta \in \mathbf{R}^{n-2}$ has entries in the interval $(0, \frac{\pi}{2})$. Notice

$$
\sin B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \operatorname{Diag} s & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \cos B = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \operatorname{Diag} c & 0 \\ 0 & 0 & 1 \end{bmatrix},
$$

where the entries $s_j, c_j \in (0, 1)$ satisfy $s_j^2 + c_j^2 = 1$ for all $j = 1, 2, 3, \ldots, n-2$, so

$$
U(\cos B) - w(\sin B) = \begin{bmatrix} -w & c_1 e_1^* & 0 \\ 0 & J(\operatorname{Diag} c) - w(\operatorname{Diag} s) & e_{n-2} \\ 0 & 0 & 0 \end{bmatrix}.
$$

A calculation shows that the matrix $G = \operatorname{Diag} g$, where

$$
g_1 = 1 \quad \text{and} \quad g_{j+1} = \frac{s_{j-1}}{c_j} g_j \quad (j = 1, 2, 3, \ldots, n-1)
$$

(setting $s_0 = 1 = c_{n-1}$) satisfies the assumptions of Theorem 8.3. Hence the numerical radius is partly smooth, relative to the manifold $\mathcal{A}$, around the matrix

$$
X = 2(\sin B)U(\cos B) = 2 \begin{bmatrix} 0 & c_1 e_1^* & 0 \\ 0 & (\operatorname{Diag} s)J(\operatorname{Diag} c) & s_{n-2} e_{n-2} \\ 0 & 0 & 0 \end{bmatrix},
$$

or more explicitly,

$$
(8.4) \qquad X = 2 \begin{bmatrix} 0 & s_0 c_1 & 0 & \cdots & 0 \\ 0 & 0 & s_1 c_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & s_{n-2} c_{n-1} \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}.
$$

23

For example, setting $s_j = \frac{1}{\sqrt{2}} = c_j$ for all $j = 1, 2, 3, \ldots, n-2$ gives

$$X = \begin{bmatrix} 0 & \sqrt{2} & & & & \\ & \cdot & 1 & & & \\ & & \cdot & \cdot & & \\ & & & \cdot & \cdot & \\ & & & & \cdot & 1 \\ & & & & & \sqrt{2} \\ & & & & & \cdot & 0 \end{bmatrix},$$

the Crabb matrix. In fact, we have the following more general result. We begin with a simple tool concerning real sequences. For notational simplicity, and with no loss of generality, we present it for infinite sequences.

Lemma 8.5 *For any strictly positive real sequence $b_1, b_2, b_3, \ldots$, there exists a decreasing sequence*

$$+\infty = t_1 > t_2 > t_3 > \cdots > 0$$

and a sequence of continuous, strictly increasing functions $h_j \colon [0, t_j) \to \mathbf{R}_+$, for all $j = 1, 2, 3, \ldots$, satisfying

$$h_j(0) = 0, \quad h_j(t_{j+1}) = 1, \quad \lim_{t \uparrow t_j} h_j(t) = +\infty,$$

$h_1(t) = b_1 t$ *for all $t \in \mathbf{R}_+$, and*

(8.6) $$h_{j+1}(t) = \frac{b_{j+1} t}{1 - h_j(t)} \quad \text{for all } t \in [0, t_{j+1}).$$

Proof We construct the numbers t_j and the functions h_j recursively. Given the continuous, strictly increasing function $h_j \colon [0, t_j) \to \mathbf{R}_+$ satisfying $h_j(0) = 0$ and $\lim_{t \uparrow t_j} h_j(t) = +\infty$ (as holds in particular when $j = 1$), by the intermediate value theorem there exists a number $t_{j+1} \in (0, t_j)$ such that $h_j(t_{j+1}) = 1$. Now the function h_{j+1} defined by equation (8.6) is continuous, strictly increasing and satisfies $h_{j+1}(0) = 0$ and $\lim_{t \uparrow t_{j+1}} h_{j+1}(t) = +\infty$, as required. $\square$

We can now prove the following result.

Theorem 8.7 (Partial smoothness at superdiagonal matrices)
At any matrix in $\mathbf{M}^n$ *(with* $n > 1$*) of the form*

$$(8.8) \qquad \begin{bmatrix} 0 & a_1 & & & & \\ & \cdot & a_2 & & & \\ & & \cdot & \cdot & & \\ & & & \cdot & \cdot & \\ & & & & \cdot & a_{n-2} \\ & & & & & \cdot & a_{n-1} \\ & & & & & & 0 \end{bmatrix}, \qquad \textit{for nonzero } a_1, a_2, \ldots, a_{n-1} \in \mathbf{C},$$

the numerical radius is partly smooth relative to the set $\mathcal{A}$ *of disk matrices.*

Proof We covered the case $n = 2$ in Section 8.1, so assume $n > 2$. We claim that any matrix of the form (8.8) is unitarily similar to a nonzero multiple of a matrix of the form (8.4). To see this, first note that, by diagonal unitary scaling we can first assume that each scalar a_j is real and strictly positive. We now seek numbers $s_j, c_j \in (0, 1)$ satisfying $s_j^2 + c_j^2 = 1$ for all $j = 1, 2, 3, \ldots, n - 2$ and such that the vector

$$\begin{bmatrix} s_0 c_1 & s_1 c_2 & s_2 c_3 & \cdots & s_{n-3} c_{n-2} & s_{n-2} c_{n-1} \end{bmatrix}$$

is a multiple of the vector

$$\begin{bmatrix} a_1 & a_2 & a_3 & \cdots & a_{n-2} & a_{n-1} \end{bmatrix}.$$

As in the discussion before equation (8.4), we set $s_0 = 1 = c_{n-1}$. Squaring componentwise, we can write this problem as solving, for $c_j \in (0, 1)$ and $t > 0$, the equations

$$c_1^2 = tb_1, \ (1 - c_j^2)c_{j+1}^2 = tb_{j+1} \ (j = 2, \ldots, n - 2), \ c_{n-1} = 1,$$

where $b_j = a_j^2$ for all j. Using the notation of the lemma, these equations are equivalent to $c_j^2 = h_j(t)$ for $j \leq n - 1$, so are solvable: the unique solution is $t = t_n$ and $c_j = \sqrt{h_j(t_n)}$ for $j \leq n - 1$.

   As we claimed, therefore, our matrix is unitarily similar to a nonzero multiple of a matrix of the form (8.4). As in the discussion of that example, the result now follows from Theorem 8.3. $\qquad \square$

We could, if desired, compute the numbers c_j recursively as described in the proof of Lemma 8.5.

9   The three-by-three case

We end with a thorough study of 3-by-3 disk matrices. In particular, we show that Theorem 7.4 has applications beyond matrices with the sparsity pattern of a Jordan block and their unitary similarity transformations.

Any disk matrix has a zero eigenvalue with algebraic multiplicity at least two, as shown in [10, p. 438]. Via Schur factorization, appropriately ordering the eigenvalues, any disk matrix in $\mathbf{M}^3$ is therefore unitarily similar to a matrix of the form

$$
(9.1) \qquad \begin{bmatrix} 0 & S \\ 0 & 0 \end{bmatrix} \quad \text{with} \quad S = 2 \begin{bmatrix} a & b \\ c & d \end{bmatrix}
$$

for numbers $a, b, c, d \in \mathbf{C}$. (The factor 2 is just for convenience, to simplify our later calculations.) Furthermore, as we show below (in Theorem 9.10), the condition

$$
(9.2) \qquad c(|a|^2 + |d|^2) = -adb^*
$$

must hold. In other words, if we define two nonconvex cones of matrices,

$$
\Xi = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} : \text{condition (9.2) holds} \right\}
$$

and

$$
\mathcal{E} = \left\{ U \begin{bmatrix} 0 & S \\ 0 & 0 \end{bmatrix} U^* : U \in \mathbf{M}^3 \text{ unitary}, \ S \in \Xi \right\},
$$

then, in the space $\mathbf{M}^3$, the set of disk matrices $\mathcal{A}$ is contained in $\mathcal{E}$. However, as we shall also see, this containment is strict.

To organize the discussion, assuming condition (9.2), exactly one of the following additional labeled conditions must hold:

$$
\begin{aligned}
a \neq 0 \ \text{ and } \ d \neq 0 & \qquad (\neq\neq) \\
a = 0 \ \text{ and } \ d \neq 0 & \qquad (=\neq) \\
a \neq 0 \ \text{ and } \ d = 0 & \qquad (\neq=) \\
a = 0 \ \text{ and } \ d = 0 \ \text{ and } \ 2|c| < |b| & \qquad (==<) \\
a = 0 \ \text{ and } \ d = 0 \ \text{ and } \ 2|c| = |b| & \qquad (===) \\
a = 0 \ \text{ and } \ d = 0 \ \text{ and } \ 2|c| > |b| & \qquad (==>)
\end{aligned}
$$

Notice that the second and third conditions both entail $c = 0$, by condition (9.2), resulting in matrices of the form (9.1) that are nilpotent.

Based on these mutually exclusive conditions, we can partition the set Ξ into corresponding disjoint subsets,

$$
(9.3) \qquad \Xi = \Xi_{\neq\neq} \cup \Xi_{=\neq} \cup \Xi_{\neq=} \cup \Xi_{==<} \cup \Xi_{===} \cup \Xi_{==>},
$$

and we can then write the set $\mathcal{E}$ as a union of corresponding subsets,

$$
\mathcal{E} = \mathcal{E}_{\neq\neq} \cup \mathcal{E}_{=\neq} \cup \mathcal{E}_{\neq=} \cup \mathcal{E}_{==<} \cup \mathcal{E}_{===} \cup \mathcal{E}_{==>}.
$$

This latter union is not, however, disjoint, as the following observation makes clear.

Proposition 9.4 *The sets $\mathcal{E}_{=\neq}$ and $\mathcal{E}_{\neq=}$ are identical and consist of the nilpotent matrices in the set $\mathcal{E}_{==<}$.*

Proof We simply observe, for example, that the matrices

$$\begin{bmatrix} 0 & 0 & b \\ 0 & 0 & d \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 & 0 & |b|^2 + |d|^2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

are unitarily similar. To see this, consider the unitary similarity transformation corresponding to $U = \begin{bmatrix} V & 0 \\ 0 & 1 \end{bmatrix}$, where the 2-by-2 unitary matrix V maps the vector $\begin{bmatrix} b \\ d \end{bmatrix}$ to the vector $\begin{bmatrix} |b|^2 + |d|^2 \\ 0 \end{bmatrix}$. $\qquad\square$

With this result in mind, we can write

$$\mathcal{E} \;\; = \;\; \mathcal{E}_{\neq\neq} \;\cup\; \mathcal{E}_{==<} \;\cup\; \mathcal{E}_{===} \;\cup\; \mathcal{E}_{==>},$$

and, as we shall see, this is indeed a partition into disjoint sets, and one well suited to understanding the set $\mathcal{A}$ of disk matrices and the numerical radius r. Specifically, we devote this section to showing the following result. Recall that Ω is the open subset of $\mathbf{M}^3$ defined in equation (5.6).

Theorem 9.5 *In the 18-dimensional space $\mathbf{M}^3$, the set $\mathcal{E}_{\neq\neq}$ is a 12-dimensional manifold. Its closure is the set of disk matrices, which can be partitioned into three disjoint sets:*

$$\mathcal{A} \;\; = \;\; \mathrm{cl}\,\mathcal{E}_{\neq\neq} \;\; = \;\; \mathcal{E}_{\neq\neq} \;\cup\; \mathcal{E}_{==<} \;\cup\; \mathcal{E}_{===}.$$

These sets satisfy

$$\mathcal{A} \cap \Omega \;\; = \;\; \mathcal{E}_{\neq\neq} \;\cup\; \mathcal{E}_{==<}.$$

The subdifferential of the numerical radius $\partial r(\cdot)$ has dimension 6 throughout $\mathcal{E}_{\neq\neq}$ and 4 throughout $\mathcal{E}_{==<}$. The numerical radius is partly smooth at a disk matrix X relative to $\mathcal{A}$ if and only if $X \in \mathcal{E}_{\neq\neq}$.

We begin the proof with some simple algebraic observations about the partition (9.3).

Proposition 9.6 *The following relationships hold:*

$$\mathrm{cl}\,\Xi_{\neq\neq} = \Xi \setminus \Xi_{==>} \quad \text{and} \quad \mathrm{cl}\,\mathcal{E}_{\neq\neq} = \mathcal{E} \setminus \mathcal{E}_{==>}.$$

Proof For any number $b \in \mathbf{C}$ and nonzero $d \in \mathbf{C}$, we have

$$\lim_{\epsilon \to 0} \begin{bmatrix} \epsilon & b \\ -\frac{db^*\epsilon}{\epsilon^2 + |d|^2} & d \end{bmatrix} \;\; = \;\; \begin{bmatrix} 0 & b \\ 0 & d \end{bmatrix},$$

27

so

(9.7)
$$\Xi_{=\neq} \subset \operatorname{cl}(\Xi_{\neq\neq}),$$

and similarly $\Xi_{\neq=} \subset \operatorname{cl}(\Xi_{\neq\neq})$. On the other hand, for any numbers $c, d \in \mathbf{C}$ satisfying $2|c| < |b|$, the real number

$$t = \frac{1}{2}\left(\sqrt{1 - \left(\frac{2|c|}{|b|}\right)^2} - 1\right)$$

satisfies

$$t^2 + t + \frac{|c|^2}{|b|^2} = 0,$$

so for $\epsilon > 0$ we have

$$\begin{bmatrix} \epsilon c & b \\ c & \epsilon bt \end{bmatrix} \in \begin{cases} \Xi_{\neq\neq} & (c \neq 0) \\ \Xi_{=\neq} & (c = 0). \end{cases}$$

This matrix converges to the matrix $\begin{bmatrix} 0 & b \\ c & 0 \end{bmatrix}$ as $\epsilon \downarrow 0$, so using property (9.7) we deduce $\Xi_{==<} \subset \operatorname{cl}(\Xi_{\neq\neq})$. Since we obviously have

(9.8)
$$\Xi_{===} \subset \operatorname{cl}(\Xi_{==<}),$$

we deduce $\operatorname{cl}(\Xi_{\neq\neq}) \supset \Xi \setminus \Xi_{==>}$. The opposite inclusion is easy: when a and d are nonzero, condition (9.2) implies $2|c| \leq |b|$, since

$$\frac{|a|^2 + |d|^2}{|ad|} \geq 2.$$

The second equation is a simple consequence, using the compactness of the set of unitary matrices. $\square$

We next introduce a simple tool.

Lemma 9.9 *Consider a matrix X unitarily similar to a matrix of the form (9.1). The characteristic polynomial of the matrix $\Phi(w^*X)$, namely $\det\big(\Phi(w^*X) - \lambda I\big)$, is*

$$-\lambda^3 + 2\lambda^2\langle w, c\rangle + (|a|^2 + |d|^2 + |b|^2)\lambda + 2\langle w, adb^* - c|b|^2\rangle.$$

Proof If X is unitarily similar to a matrix E in the form (9.1), then the characteristic polynomial of $\Phi(w^*X)$ is identical with that of $\Phi(w^*E)$, namely

$$\det\big(\Phi(w^*E) - \lambda I\big) = \det \begin{bmatrix} -\lambda & w^*a & w^*b \\ wa^* & wc^* + w^*c - \lambda & w^*d \\ wb^* & wd^* & -\lambda \end{bmatrix},$$

and the result follows. $\square$

We can now characterize disk matrices in $\mathbf{M}^3$.

28

Theorem 9.10 *The set of disk matrices in M^3 is $\mathcal{E} \setminus \mathcal{E}_{==>}$.*

Proof We need to show that a matrix $X \in \mathbf{M}^3$ is a disk matrix if and only if it is unitarily similar to a matrix of the form $\left[\begin{smallmatrix} 0 & S \\ 0 & 0 \end{smallmatrix}\right]$ for a matrix $S \in \Xi \setminus \Xi_{==>}$. To this end, consider any X with field of values a disk of radius $r \geq 0$. As we noted, X is unitarily similar to a matrix of the form (9.1), so by Lemma 9.9 the characteristic polynomial of $\Phi(w^*X)$ is

$$-\lambda^3 + 2\lambda^2\langle w, c\rangle + (|a|^2 + |d|^2 + |b|^2)\lambda + 2\langle w, adb^* - c|b|^2\rangle.$$

By Proposition 5.3, for all $w \in \mathbf{T}$ we have

$$-r^3 + 2r^2\langle w, c\rangle + (|a|^2 + |d|^2 + |b|^2)r + 2\langle w, adb^* - c|b|^2\rangle \;=\; 0.$$

By considering $w = \pm 1$ and adding we deduce

$$r = \sqrt{|a|^2 + |d|^2 + |b|^2}.$$

Substituting into the previous equation gives

$$\langle w, c(|a|^2 + |d|^2) + adb^*\rangle \;=\; 0$$

for all $w \in \mathbf{T}$, and equation (9.2) follows.

If in addition we have $S \in \Xi_{==>}$, then $r = |b|$. Our previous calculation shows

$$\begin{aligned}
\det\big(\Phi(w^*E) &- \lambda I\big) \\
&= -\lambda^3 + 2\lambda^2\langle w, c\rangle + |b|^2\lambda - 2|b|^2\langle w, c\rangle \\
&= -(\lambda - 2\langle w, c\rangle)\big(\lambda^2 - |b|^2\big),
\end{aligned}$$

so the matrix $\Phi(w^*X)$ has eigenvalues $2\langle w, c\rangle$ and $\pm|b|$. Setting $w = \frac{1}{|c|}c$ now gives $r = \max(2|c|, |b|) = 2|c|$, a contradiction.

Conversely, consider any matrix E of the form (9.1), where $S \in \Xi$ but $S \notin \Xi_{==>}$. We want to prove

$$\lambda_{\max}\big(\Phi(w^*E)\big) = \sqrt{|a|^2 + |d|^2 + |b|^2} \quad \text{for all } w \in \mathbf{T},$$

which will imply that E, along with any matrix that is unitarily similar to E, is a disk matrix. Lemma 9.9 shows

$$\begin{aligned}
\det\big(\Phi(w^*E) &- \lambda I\big) \\
&= -\lambda^3 + 2\lambda^2\langle w, c\rangle + (|a|^2 + |d|^2 + |b|^2)\lambda - 2(|a|^2 + |d|^2 + |b|^2)\langle w, c\rangle \\
&= -(\lambda - 2\langle w, c\rangle)\big(\lambda^2 - (|a|^2 + |d|^2 + |b|^2)\big),
\end{aligned}$$

29

so the matrix $\Phi(w^*E)$ has eigenvalues $2\langle w, c \rangle$ and $\pm\sqrt{|a|^2 + |d|^2 + |b|^2}$. The result now follows immediately in the case when $a = 0 = d$, since then by assumption $2|c| \leq |b|$. It also follows if either a or d is nonzero, since equation (9.2) then implies

$$2|c| < \sqrt{|a|^2 + |d|^2 + |b|^2}$$

To see this, note the equivalent form

$$2|c|(|a|^2 + |d|^2) < (|a|^2 + |d|^2)\sqrt{|a|^2 + |d|^2 + |b|^2}$$

or equivalently

$$2|a||d||b| < (|a|^2 + |d|^2)\sqrt{|a|^2 + |d|^2 + |b|^2}.$$

Squaring gives

$$4|a|^2|d|^2|b|^2 < (|a|^2 + |d|^2)^2(|a|^2 + |d|^2 + |b|^2)$$

or equivalently the valid inequality

$$0 < (|a|^2 - |d|^2)^2|b|^2 + (|a|^2 + |d|^2)^3.$$

This completes the proof of the converse direction. $\qquad\square$

Examining the proof, we have in particular shown the following result. Notice that, having already discarded the subset $\Xi_{==>}$, this reveals the special role of the subset $\Xi_{===}$.

Lemma 9.11 *A matrix E of the form (9.1) is a disk matrix if and only if the matrix S lies in the set $\Xi \setminus \Xi_{==>}$. In that case, E has numerical radius*

$$\lambda_{\max}\big(\Phi(w^*E)\big) = \sqrt{|a|^2 + |d|^2 + |b|^2} \ \ \text{for all } w \in \mathbf{T}.$$

Furthermore this largest eigenvalue is simple for all $w \in \mathbf{T}$ if and only if $S \notin \Xi_{===}$.

Discarding the set $\Xi_{===}$ rules out the case of the zero matrix, so by scaling we can focus our attention on matrices with field of values the unit disk. In this context, we next further discard the set $\Xi_{==<}$.

Lemma 9.12 *A matrix E of the form (9.1) has field of values the unit disk if and only if the matrix S lies in the set $\Xi \setminus \Xi_{==>}$ and $|a|^2 + |d|^2 + |b|^2 = 1$. Assume this is the case. Then*

$$\lambda_{\max}\big(\Phi(w^*E)\big) = 1 \ \ \text{for all } w \in \mathbf{T}.$$

Furthermore, this largest eigenvalue is simple for all $w \in \mathbf{T}$ if and only if we have $S \notin \Xi_{===}$. If we further assume that a and d are not both zero (so $S \notin \Xi_{==<}$), then a corresponding unit eigenvector is

$$\frac{1}{\sqrt{2(1 - |b|^2)(1 - 2\langle w, c \rangle)}} \begin{bmatrix} w^*a + bd^* \\ 1 - |b|^2 \\ b^*a + wd^* \end{bmatrix}.$$

Proof By our previous calculations we know, for all $w \in \mathbf{T}$, that the matrix

$$\Phi(w^*E) - I = \begin{bmatrix} -1 & w^*a & w^*b \\ wa^* & wc^* + w^*c - 1 & w^*d \\ wb^* & wd^* & -1 \end{bmatrix}$$

has nullity one. It is easy to verify

$$\begin{bmatrix} -1 & w^*a & w^*b \\ wb^* & wd^* & -1 \end{bmatrix} \begin{bmatrix} w^*a + bd^* \\ 1 - |b|^2 \\ b^*a + wd^* \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

so the given vector is indeed an eigenvector, and the square of its norm is

$$\frac{|a|^2 + |b|^2|d|^2 + 2\langle w, adb^* \rangle + (1 - |b|^2)^2 + |a|^2|b|^2 + |d|^2 + 2\langle w, adb^* \rangle}{2(1 - |b|^2)(1 - 2\langle w, c \rangle)} = 1$$

as required. $\qquad\square$

We are now ready to consider partial smoothness.

Theorem 9.13 *Consider any matrix E of the form (9.1), for a matrix S in the set $\Xi_{\neq\neq}$. Then E belongs to the set $\mathcal{A}$ of disk matrices. Around E, furthermore, $\mathcal{A}$ is a 12-dimensional analytic manifold in the 18-dimensional space $\mathbf{M}^3$, with normal space at E given by*

$$N_{\mathcal{A}}(E) = E^{\perp} \cap G\Big(\mathrm{span}\{F(w) : w \in \mathbf{T}\}\Big)G^*,$$

where G is the invertible matrix

$$G = \begin{bmatrix} a & bd^* & 0 \\ 0 & 1 - |b|^2 & 0 \\ 0 & ab^* & d^* \end{bmatrix}.$$

The numerical radius is partly smooth at E relative to $\mathcal{A}$, and its subdifferential there, $\partial r(E)$, has dimension 6.

Proof As in the proof of Theorem 7.4, we lose no generality in first scaling the matrix E to ensure $|a|^2 + |d|^2 + |b|^2 = 1$. Then, by Lemma 9.12, the eigenvalue

$$\lambda_{\max}\big(\Phi(w^*E)\big) = 1$$

is simple for all scalars $w \in \mathbf{T}$, with a corresponding eigenvector

$$Gf(w) = \begin{bmatrix} aw + bd^*w^2 \\ (1 - |b|^2)w^2 \\ b^*aw^2 + d^*w^3 \end{bmatrix}.$$

The result now follows by Theorem 7.4. $\qquad\square$

31

Corollary 9.14 *The semi-algebraic set $\mathcal{A}$ of all disk matrices in $\mathbf{M}^3$ has dimension 12.*

Proof By Proposition 9.6 and Theorem 9.10, any disk matrix is a limit of matrices X unitarily similar to those of the form considered in Theorem 9.13. That result shows that the set of such matrices X is a semi-algebraic manifold of dimension 12, so its closure also has dimension 12. □

By contrast, we have the following result.

Proposition 9.15 *Any matrix $X \in \mathcal{E}_{==<}$ is a disk matrix: in fact $X \in \mathcal{A} \cap \Omega$. However, the subdifferential of the numerical radius $\partial r(X)$ has dimension 4. Consequently, the numerical radius is not partly smooth at X relative to any subset of $\mathcal{A}$.*

Proof After a unitary similarity transformation, we can suppose

$$
X \;=\; 2 \begin{bmatrix} 0 & 0 & b \\ 0 & c & 0 \\ 0 & 0 & 0 \end{bmatrix},
$$

with $2|c| < |b|$. The claim $X \in \mathcal{A} \cap \Omega$ follows from Lemma 9.12. By Lemma 9.11, for all scalars $w \in \mathbf{T}$, the eigenvalue

$$
\lambda_{\max}\big(\Phi(w^*X)\big) = |b|
$$

is simple, and a quick calculation shows that a corresponding unit eigenvector is

$$
\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 & uw^* \end{bmatrix}^*
$$

where $u = \frac{b}{|b|}$. Now Theorem 5.8 implies that the numerical radius has subdifferential at X given by

$$
\partial r(X) \;=\; \frac{1}{2}\mathrm{conv}\left\{ \begin{bmatrix} w & 0 & u \\ 0 & 0 & 0 \\ u^*w^2 & 0 & w \end{bmatrix} : w \in \mathbf{T} \right\}.
$$

We can rewrite the matrix in this equation as the image of the vector $[1 \ w \ w^2]$ under a one-to-one linear map, so by Corollary 6.3 we see that $\partial r(X)$ has dimension 4. Finally, by definition, any manifold with respect to which the numerical radius is partly smooth at X must have dimension $18 - 4 = 14$, and hence cannot be a subset of the 12-dimensional set $\mathcal{A}$. □

We can summarize some of our observations in the following refinement of Theorem 9.10, giving the final ingredient in the proof of Theorem 9.5.

Corollary 9.16 *For any matrix X in the set of disk matrices $\mathcal{A}$ in $\mathbf{M}^3$, the numerical radius is partly smooth at X relative to $\mathcal{A}$ if and only if $X \in \mathcal{E}_{\neq\neq}$.*

Proof We can suppose that X is unitarily similar to a matrix $E = \left[\begin{smallmatrix} 0 & 2S \\ 0 & 0 \end{smallmatrix}\right]$ for a matrix $S \in \Xi \setminus \Xi_{==>}$. We claim that the numerical radius is partly smooth at E relative to $\mathcal{A}$ if and only if $S \in \Xi_{\neq\neq}$. In the light of Theorem 9.13 and Proposition 9.15, the only case left to consider is $S \in \Xi_{===}$. In the case $S \in \Xi_{\neq\neq}$ we know $\dim \partial r(E) = 6$, whereas in the case $S \in \Xi_{==<}$ we have $\dim \partial r(E) = 4$. But when $S \in \Xi_{===}$, our matrix can be written both as a limit of matrices $E \in \mathcal{A}$ with $S \in \Xi_{\neq\neq}$ or alternatively with $S \in \Xi_{==<}$, by Proposition 9.6 and equation (9.8), so $\dim \partial r$ is not continuous there relative to $\mathcal{A}$, contradicting partial smoothness. $\square$

It is interesting to consider the impact of these results on optimization problems involving the numerical radius r on $\mathbf{M}^3$. For example, by Theorems 3.4 and 9.13, the prox operator prox_r commonly maps a random matrix Y to a disk matrix, which is therefore unitarily similar to a matrix $E = \left[\begin{smallmatrix} 0 & 2S \\ 0 & 0 \end{smallmatrix}\right]$ for a matrix $S \in \Xi \setminus \Xi_{==>}$. In particular, the case $S \in \Xi_{\neq\neq}$ occurs with positive probability, and this property is stable under perturbations to Y.

Worth noting is that the case $S \in \Xi_{===}$ can also occur stably with positive probability. For example, the prox operator prox_r maps any small matrix to the zero matrix.

On the other hand, the case $S \in \Xi_{==<}$ covered by Proposition 9.15 (as well as the cases $S \in \Xi_{=\neq} \cup \Xi_{\neq=}$) do not display this same stable behavior. For almost all matrices Y, the numerical radius, being semi-algebraic, will indeed be partly smooth around $\mathrm{prox}_r(Y)$ relative to some manifold $\mathcal{M}$ (see [6]). However, even if $\mathrm{prox}_r(Y)$ were a disk matrix of one of these three forms, $\mathcal{M}$ could not be a subset of $\mathcal{A}$ locally, so certain small perturbations to Y would produce matrices other than disk matrices.

The Crouzeix ratio in the 3-by-3 case

It is interesting to compare how two 3×3 families given in [10] fit into the discussion of this section. Let

$$Y(\xi) = \begin{bmatrix} 0 & 0 & 2 \\ 0 & \xi & 0 \\ 0 & 0 & 0 \end{bmatrix}, \ \xi \in \mathbf{C}, \ |\xi| \le 1$$

and

$$Z(\phi,\psi) = e^{i\psi} \begin{bmatrix} 0 & \sqrt{2}\cos\phi & 2\sin\phi \\ 0 & -\sin\phi & \sqrt{2}\cos\phi \\ 0 & 0 & 0 \end{bmatrix}, \ \phi \in \left[0, \frac{\pi}{2}\right], \ \psi \in \mathbf{R}.$$

33

Crouzeix proved in [10, Lemma 10.2] that these matrices, along with their unitary similarity transformations, are the only 3×3 matrices X with field of values the unit disk and

$$\inf \left\{ \frac{\max\{|p(\lambda)| : |\lambda| \leq 1\}}{\|p(X)\|_2} : p \text{ is a polynomial of any degree} \right\} = \frac{1}{2}.$$

Note that the ratio in (2.1) is a special case of the one in the equation above, with $p(\lambda) = \lambda^{n-1}$. Indeed, since permuting the rows and columns of $Y(\xi)$ results in a block diagonal matrix with one block equal to the scaled 2×2 Jordan block discussed in Section 4, and a second block equal to the scalar ξ with $|\xi| \leq 1$, when $X = Y(\xi)$ the infimum in the equation above is attained by $p(\lambda) = \lambda$. When $\phi = \frac{\pi}{2}$, the matrix $Z(\phi, \psi)$ can also, by permuting rows and columns, be transformed to essentially the same block diagonal form, while when $\phi = 0$, $Z(\phi, \psi)$ is a unimodular scalar times the scaled 3×3 Jordan block discussed in Section 4, and the infimum is attained by $p(\lambda) = \lambda^2$. In general, however, for $X = Z(\phi, \psi)$ the infimum is not attained (although it is if p is allowed to be any analytic function).

We can apply the results in this section directly to the matrices $Y(\xi)$ and $Z(\phi, \psi)$. In particular, the numerical radius is partly smooth relative to the set of disk matrices at the matrix $Z(\phi, \psi)$ if and only if $\phi \neq \frac{\pi}{2}$, and is not so at the matrix $Y(\xi)$.

10   Concluding Remarks

In this paper, we have analyzed the partly smooth geometry of the numerical radius around disk matrices — those whose fields of values are disks in the complex plane centered at the origin. We have argued that such matrices arise naturally as solutions of optimization problems involving the numerical radius, supported by both numerical experiments and subdifferential analysis. Even in the case of matrices of order three, the details are surprisingly subtle. Our development blends matrix and variational analysis, but makes no use of semidefinite programming characterizations of the numerical radius like Theorem 1.2: the implication of this characterization for partial smoothness is an interesting open question.

References

[1] A.G. Aksoy and G. Lewicki. Best approximation in numerical radius. *Numer. Funct. Anal. Optim.*, 32(6):593–609, 2011.

[2] T. Ando. Structure of operators with numerical radius one. *Acta Sci. Math. (Szeged)*, 34:11–15, 1973.

[3] C.A. Berger. A strange dilation theorem. *Notices Amer. Math. Soc.*, 12:590, 1965.

[4] C.A. Berger and J.G. Stampfli. Mapping theorems for the numerical range. *Amer. J. Math.*, 89:1047–1055, 1967.

[5] V. Blondel, M. Gevers, and A. Lindquist. Survey on the state of systems and control. *European Journal of Control*, 1:5–23, 1995.

[6] J. Bolte, A. Daniilidis, and A.S. Lewis. Generic optimality conditions for semi-algebraic convex programs. *Math. Oper. Res.*, 36:55–70, 2011.

[7] Sheung Hun Cheng and N.J. Higham. The nearest definite pair for the Hermitian generalized eigenvalue problem. *Linear Algebra and its Applications*, 302-303:63–76, 1999.

[8] M.J. Crabb. The powers of an operator of numerical radius one. *Michigan Math. J.*, 18:253–256, 1971.

[9] M. Crouzeix. Bounds for analytical functions of matrices. *Integral Equations Operator Theory*, 48(4):461–477, 2004.

[10] M. Crouzeix. Some constants related to numerical ranges. *SIAM J. Matrix Anal. Appl.*, 37(1):420–442, 2016.

[11] M. Crouzeix and C. Palencia. The numerical range is a $(1 + \sqrt{2})$-spectral set. *SIAM J. Matrix Anal. Appl.*, 38(2):649–655, 2017.

[12] T.A. Driscoll, N. Hale, and L.N. Trefethen. *Chebfun Guide*. Pafnuty Publications, Oxford, 2014.

[13] D. Drusvyatskiy and A.S. Lewis. Optimality, identifiability, and sensitivity. *Math. Program.*, 147:467–498, 2014.

[14] F.R. Gantmacher. *Matrix Theory*, volume 1. Chelsea Publishing Company, 1960.

[15] M. Goldberg and E. Tadmor. On the numerical radius and its applications. *Linear Algebra Appl.*, 42:263–284, 1982.

[16] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1, March 2014. http://cvxr.com/cvx.

[17] A. Greenbaum, A.S. Lewis, and M.L. Overton. Variational analysis of the Crouzeix ratio. *Math. Program.*, 164:229–243, 2017.

[18] A. Greenbaum and M.L. Overton. Numerical investigation of Crouzeix's conjecture. *Linear Algebra and its Applications*, 542:225–245, 2018.

[19] K.E. Gustafson and D.K.M. Rao. *Numerical Range*. Springer-Verlag, New York, 1997.

[20] R.A. Horn and C.R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, second edition, 2013.

[21] R. Kippenhahn. Uber den Wertevorrat einer Matrix. *Math. Nachr.*, 6:193–228, 1951.

[22] P.D. Lax. *Linear Algebra and its Applications*. Wiley-Interscience, Hoboken, NJ, second edition, 2007.

[23] P.D. Lax and B. Wendroff. Difference schemes for hyperbolic equations with high order of accuracy. *Comm. Pure Appl. Math.*, 17:381–398, 1964.

[24] A.S. Lewis. Active sets, nonsmoothness, and sensitivity. *SIAM J. Optim.*, 13(3):702–725, 2002.

[25] A.S. Lewis and S. Zhang. Partial smoothness, tilt stability, and generalized Hessians. *SIAM J. Optim.*, 23(1):74–94, 2013.

[26] C.-K. Li. Numerical ranges and numerical radii. people.wm.edu/∼cklixx/wonra.html.

[27] C.-K. Li. Numerical range. In L. Hogben, editor, *Handbook of Numerical Analysis*, pages 18–1 – 18–12. Chapman and Hall, 2007.

[28] G. Lumer. Semi-inner-product spaces. *Trans. Amer. Math. Soc.*, 100:29–43, 1961.

[29] R. Mathias. Matrix completions, norms and Hadamard products. *Proc. Amer. Math. Soc.*, 117(4):905–918, 1993.

[30] 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.

[31] C. Pearcy. An elementary proof of the power inequality for the numerical radius. *Michigan Math. J.*, 13:289–291, 1966.

[32] M.J.D. Powell. *Approximation Theory and Methods*. Cambridge University Press, Cambridge-New York, 1981.

[33] R.T. Rockafellar and R.J-B. Wets. *Variational Analysis*. Grundlehren der mathematischen Wissenschaften, Vol 317, Springer, Berlin, 1998.

[34] L.N. Trefethen and M. Embree. *Spectra and Pseudospectra*. Princeton University Press, Princeton, NJ, 2005.

[35] N.-K. Tsing, M.K.H. Fan, and E.I. Verriest. On analyticity of functions involving eigenvalues. *Linear Algebra and its Applications*, 207:159–180, 1994.

[36] F. Uhlig. Geometric computation of the numerical radius of a matrix. *Numerical Algorithms*, 52:335–353, 2009.

[37] X.A. Wang. Grassmannian, central projection, and output feedback pole assignment of linear systems. *IEEE Transactions on Automatic Control*, 41:786–794, 1996.

[38] J. C. Willems. Generic eigenvalue assignability by real memoryless output feedback made simple. In A. Paulraj, V. Roychowdhury, and C.D. Schaper, editors, *Communications, Computation, Control, and Signal Processing: a tribute to Thomas Kailath*, pages 343–354. Springer US, Boston, MA, 1997.

[39] W.M. Wonham. *Linear Multivariable Control*. Springer-Verlag, New York, third edition, 1985.