p-adic differential equations
18.787, Kiran S. Kedlaya, MIT, fall 2007
A little p-adic numerical analysis

In this unit, we introduce some concepts from numerical analysis, and their nonarchimedean analogues.

1 Singular values

Until further notice, let A be an $n \times n$ matrix over $\mathbb{C}$. One set of numerical invariants we can attach to A is the list of eigenvalues $\lambda_1, \dots, \lambda_n$, which we sort so that $|\lambda_1| \geq \cdots \geq |\lambda_n|$.

A second set, which is more useful for numerical analysis, is the singular values. Let A^* denote the conjugate transpose (or *Hermitian transpose*) of A. The matrix A^*A is real symmetric, so has nonnegative real eigenvalues. The square roots of these eigenvalues comprise the *singular values* of A; we denote them $\sigma_1, \dots, \sigma_n$ with $\sigma_1 \geq \cdots \geq \sigma_n$. These are not invariant under conjugation, but they are invariant under multiplying A on either side by a unitary matrix.

Let $\mathrm{Diag}(\sigma_1, \dots, \sigma_n)$ denote the $n \times n$ diagonal matrix D with $D_{ii} = \sigma_i$ for $i = 1, \dots, n$.

Theorem 1 (Singular value decomposition)**.** *There exist unitary $n \times n$ matrices U, V such that $UAV = \mathrm{Diag}(\sigma_1, \dots, \sigma_n)$.*

Proof. This is equivalent to showing that there is an orthonormal basis of $\mathbb{C}^n$ which remains orthogonal upon applying A. To construct it, start with a vector $v \in \mathbb{C}^n$ maximizing $|Av|/|v|$, then show that for any $w \in \mathbb{C}^n$ orthogonal to v, Aw is also orthogonal to Av. For further details, see references in the notes. □

Corollary 2. *The singular values of A^{-1} are $\sigma_n^{-1}, \dots, \sigma_1^{-1}$.*

From the singular value decomposition, we may infer a convenient interpretation of σ_i.

Corollary 3. *The number σ_i is the largest value of λ for which the following holds: for any i-dimensional subspace V of $\mathbb{C}^n$, there exists $v \in V$ nonzero such that $|Av| \leq \lambda |v|$.*

Proof. Theorem 1 provides an orthonormal basis $v_1, \dots, v_n$ of V such that $Av_1, \dots, Av_n$ is again orthogonal, and $|Av_i| = \sigma_i |v_i|$ for $i = 1, \dots, n$. Let W be the span of $v_i, \dots, v_n$; then for any i-dimensional subspace V of $\mathbb{C}^n$, $V \cap W$ is nonempty, and any $v \in V \cap W$ satisfies $|Av| \leq \sigma_i |v|$. On the other hand, if we take V to be the span of $v_1, \dots, v_i$, then we have $|Av| \geq \sigma_i |v|$ for all $v \in V$. This proves the claim. □

The relationship between the singular values and the eigenvalues is controlled by the following inequality of Weyl [Wey49]. For a vast generalization, see Theorem 26.

Theorem 4 (Weyl)**.** *We have*

$$\sigma_1 \cdots \sigma_i \geq |\lambda_1 \cdots \lambda_i| \qquad (i = 1, \dots, n),$$

with equality for $i = n$.

Proof. The equality for $i = n$ holds because $\det(A^*A) = |\det(A)|^2$. We check the inequality first for $i = 1$. Note that if we equip $\mathbb{C}^n$ with the L_2 norm, i.e.,

$$|(z_1, \dots, z_n)| = (|z_1|^2 + \cdots + |z_n|^2)^{1/2},$$

then σ_1 is the operator norm of A, that is,

$$\sigma_1 = \sup_{v \in \mathbb{C}^n - \{0\}} \{|Av|/|v|\}.$$

Since there exists $v \in \mathbb{C}^n - \{0\}$ with $Av = \lambda_1 v$, we deduce that $\sigma_1 \geq |\lambda_1|$.

For the general case, we pass from $\mathbb{C}^n$ to its i-th exterior power $\wedge^i \mathbb{C}^n$, on which A also acts. The maximum norm of an eigenvalue of this action is $|\lambda_1 \cdots \lambda_i|$, and the operator norm is $\sigma_1 \cdots \sigma_i$. Thus the previous inequality gives what we want. □

We mention in passing the following converse of Theorem 4, due to Horn [Hor54, Theorem 4].

Theorem 5. *For $\lambda_1, \dots, \lambda_n \in \mathbb{C}$ and $\sigma_1, \dots, \sigma_n \in \mathbb{R}_{\geq 0}$ satisfying*

$$\sigma_1 \cdots \sigma_i \geq |\lambda_1 \cdots \lambda_i| \qquad (i = 1, \dots, n),$$

with equality for $i = n$, there exist an $n \times n$ matrix A over $\mathbb{C}$ with singular values $\sigma_1, \dots, \sigma_n$ and eigenvalues $\lambda_1, \dots, \lambda_n$.

Equality in Weyl's theorem at an intermediate stage has a structural meaning.

Theorem 6. *Suppose that for some $i \in \{1, \dots, n-1\}$ we have*

$$\sigma_i > \sigma_{i+1}, \qquad |\lambda_i| > |\lambda_{i+1}|,$$
$$\sigma_1 \cdots \sigma_i = |\lambda_1 \cdots \lambda_i|.$$

Then there exists a unitary matrix U such that $U^{-1}AU$ is block diagonal, with the first block accounting for the first i singular values and eigenvalues, and the second block accounting for the others.

Proof. Let $v_1, \dots, v_n$ be a basis of $\mathbb{C}^n$ such that $v_1, \dots, v_i$ span the generalized eigenspaces with eigenvalues $\lambda_1, \dots, \lambda_i$. and $v_{i+1}, \dots, v_n$ span the generalized eigenspaces with eigenvales $\lambda_{i+1}, \dots, \lambda_n$. Apply the singular value decomposition to construct an orthonormal basis $w_1, \dots, w_n$ such that $Aw_1, \dots, Aw_n$ are also orthogonal and $|Aw_i| = \sigma_i |w_i|$.

Since $\sigma_i > \sigma_{i+1}$, the only vectors $v \in \wedge^i \mathbb{C}^n$ for which $|Av|/|v|$ achieves its maximum value $\sigma_1 \cdots \sigma_i$ are the nonzero multiples of $w_1 \wedge \cdots \wedge w_i$. However, this is also true for $v_1 \wedge \cdots \wedge v_i$. We conclude that $w_1, \dots, w_i$ span V; this implies that the orthogonal complement of V is spanned by $w_{i+1}, \dots, w_n$, and so is also preserved by A. This yields the desired result. □

Theorem 7. *The following are equivalent.*

(a) There exists a unitary matrix U such that $U^{-1}AU$ is diagonal.

(b) The matrix A is normal, i.e., $A^*A = AA^*$.

(c) The eigenvalues $\lambda_1, \dots, \lambda_n$ and singular values $\sigma_1, \dots, \sigma_n$ of A satisfy $|\lambda_i| = \sigma_i$ for $i = 1, \dots, n$.

Proof. It is clear that (a) implies both (b) and (c). Given (b), we can perform a joint eigenspace decomposition for A and A^*. On any common generalized eigenspace, A has some eigenvalue λ, A^* has eigenvalue $\overline{\lambda}$, and so A^*A has eigenvalue $|\lambda|^2$. This implies (c).

Given (c), Theorem 6 implies that A can be conjugated by a unitary matrix into a block diagonal matrix in which each block has a single eigenvalue and a single singular value, which coincide. Let B be such a block, with eigenvalue λ, corresponding to a subspace V of $\mathbb{C}^n$. If the common singular value is 0, then $B = 0$. Otherwise, $\lambda \neq 0$ and λ^{-1} is unitary. Hence given orthogonal eigenvectors $v_1, \dots, v_i \in V$ of B, the orthogonal complement in V of their span is preserved by B, so is either zero or contains another eigenvector v_{i+1}. This shows that B is diagonalizable, and thus is itself a scalar matrix. (One can also argue this last step using compactness of the unitary group.) ☐

In general, we can almost conjugate any matrix into a normal matrix.

Lemma 8. *For any $\eta > 1$, we can choose $U \in \mathrm{GL}_n(\mathbb{C})$ such that for $i = 1, \dots, n$, the i-th singular value of $U^{-1}AU$ is at most $\eta|\lambda_i|$. If A is semisimple, we can also take $\eta = 1$.*

Proof. Put A in Jordan normal form, then rescale so that for each eigenvalue λ, the superdiagonal terms have absolute value at most $(|\eta| - 1)|\lambda|$, and all other terms are zero. ☐

2 Perturbations (archimedean case)

Another inequality of Weyl [Wey12] shows that the singular values do not change much under a small (additive) perturbation.

Theorem 9 (Weyl)**.** *Let B be an $n \times n$ matrix, and let $\sigma'_1, \dots, \sigma'_n$ be the singular values of $A + B$. Then*

$$|\sigma'_i - \sigma_i| \leq |B| \qquad (i = 1, \dots, n).$$

It is more complicated to describe what happens to the eigenvalues under a small additive perturbation, but it is not so difficult to quantify what happens to the characteristic polynomial, at least in a crude fashion.

Theorem 10. *Let B be an $n \times n$ matrix such that $|B| < \sigma_j$ for some $j \in \{1, \dots, n\}$. Let $P(T) = T^n + \sum_{i=0}^{n-1} P_i T^i$ and $Q(T) = T^n + \sum_{i=0}^{n-1} Q_i T^i$ be the characteristic polynomials of A and $A + B$. Then*

$$|P_{n-i} - Q_{n-i}| \leq \left| 2^i \binom{n}{i} \right| \sigma_1 \cdots \sigma_{i-1} |B| \qquad (i = 1, \dots, j).$$

The superfluous enclosure of the integer $2^n\binom{n}{i}$ in absolute value signs is quite deliberate; it will be relevant in the nonarchimedean setting.

Proof. First consider the case $i = j = n$. By continuity, we may assume that $\det(A) \neq 0$. Write

$$\begin{aligned}\det(A+B) &= \det(A)\det(I_n + A^{-1}B) \\ &= \det(A)(1 - R_{n-1} + \cdots \pm R_0),\end{aligned}$$

where $T^n + \sum_{i=0}^{n-1} R_i T^i$ is the characteristic polynomial of $A^{-1}B$. From the expansion of R_{n-i} as a sum of $\binom{n}{i}$ minors of size i, we have $|R_{n-i}| < \binom{n}{i}|A^{-1}B|^i$. Since $|A^{-1}| = \sigma_n^{-1}$, we have $|A^{-1}B| < 1$; we may thus write

$$|\det(A+B) - \det(A)| \leq |2^n||\det(A)||A^{-1}B| = |2^n|\sigma_1 \cdots \sigma_{n-1}|B|.$$

For the general case, write the coefficient of T^{n-i} in the characteristic polynomial of a matrix as the sum of $\binom{n}{i}$ minors of size i, then apply the previous case to each of these. □

We also need to consider multiplicative perturbations.

Proposition 11. *Let $B \in \mathrm{GL}_n(\mathbb{C})$ satisfy $|B| \leq \eta$. Let $\sigma'_1, \dots, \sigma'_n$ be the singular values of AB. Then*

$$\sigma'_i \leq \eta\sigma_i \qquad (i = 1, \dots, n).$$

(The analogous result holds with BA replaced by AB, since transposal does not change singular values.)

Proof. We use the interpretation of singular values given by Corollary 3. Choose an i-dimensional subspace V of $\mathbb{C}^n$ such that $|BAv| \geq \sigma'_i|v|$ for all $v \in V$. Then choose $v \in V$ nonzero such that $|Av| \leq \sigma_i|v|$. We have

$$\sigma'_i|v| \leq |BAv| \leq |B||Av| \leq \sigma_i|B||v|,$$

proving the claim. □

Proposition 12. *Let $\sigma_{k,1}, \dots, \sigma_{k,n}$ be the singular values of A^k. Then*

$$\lim_{k\to\infty} \sigma_{k,i}^{1/k} = |\lambda_i| \qquad (i = 1, \dots, n).$$

Proof. Pick $\eta > 1$, and choose U as in Lemma 8; that is, U is upper-triangular, and each block of eigenvalue λ has some scalar c of norm at most $(|\eta| - 1)|\lambda|$. Let U be the matrix effecting the resulting conjugation.

In a block with eigenvalue λ, the singular values of the k-th power are bounded below by $|\lambda|^k$ and above by $\eta^k|\lambda|^k$. Consequently, we may apply Proposition 11 to deduce that

$$|\lambda_i|^k|U||U^{-1}| \leq \sigma_{k,i} \leq \eta^k|\lambda_i|^k|U||U^{-1}|.$$

Taking k-th roots and then taking $k \to \infty$, we deduce

$$|\lambda_i| \leq \liminf_{k\to\infty} \sigma_{k,i}^{1/k}, \qquad \limsup_{k\to\infty} \sigma_{k,i}^{1/k} \leq \eta|\lambda_i|.$$

Since $\eta > 1$ was arbitrary, we deduce the desired result. □

3 Hodge and Newton polygons

We now pass to nonarchimedean analogues. For the rest of this unit, let F be a nonarchimedean field, and let A be an $n \times n$ matrix over F.

The *Hodge polygon* of A is the polygon starting from $(-n, 0)$ whose slopes $s_1, \dots, s_n$ have the property that for $i = 1, \dots, n$, $s_1 + \cdots + s_i$ is the minimum valuation of an $i \times i$ minor of A. These slopes may be more familiar as the *elementary divisors* of A; see the notes for an explanation of the term "Hodge polygon". They are also sometimes called the *invariant factors* of A.

We will refer to $\sigma_1, \dots, \sigma_n = e^{-s_1}, \dots, e^{-s_n}$ as the *singular values* of A; these are invariant under multiplication on either side by a matrix in $\mathrm{GL}_n(\mathfrak{o}_F)$. One has the relation

$$\sigma_1 = |A|,$$

but this time taking the operator norm defined by the supremum norm on F^n.

We also have an analogue of the singular value decomposition, but only when F is spherically complete.

Theorem 13 (Smith normal form)**.** *There exist $U, V \in \mathrm{GL}_n(\mathfrak{o}_F)$ such that UAV is a diagonal matrix whose entries have norms $\sigma_1, \dots, \sigma_n$.*

Proof. Exercise. □

Corollary 14. *The slopes $s_1, \dots, s_n$ of the Hodge polygon satisfy $s_1 \leq \cdots \leq s_n$.*

Proof. We may replace F by its spherical completion, and then the i-th slope s_i is evidently the i-th smallest valuation of a diagonal entry of the Smith normal form. □

Corollary 15. *The number σ_i is the largest value of λ for which the following holds: for any i-dimensional subspace V of F^n, there exists $v \in V$ nonzero such that $|Av| \leq \lambda |v|$.*

The *Newton polygon* of A is simply the Newton polygon of its characteristic polynomial. That is, for $\lambda_1, \dots, \lambda_n$ the eigenvalues of A in some algebraic extension of F, sorted with $|\lambda_1| \geq \cdots \geq |\lambda_n|$, the Newton polygon has slopes $-\log|\lambda|_1, \dots, -\log|\lambda_n|$ in that order. Since the characteristic polynomial is invariant under conjugation by an element of $\mathrm{GL}_n(F)$, so is the Newton polygon.

The archimedean analogue of Weyl's inequality is the following.

Theorem 16 (Newton above Hodge)**.** *We have*

$$\sigma_1 \cdots \sigma_i \geq |\lambda_1 \cdots \lambda_i| \qquad (i = 1, \dots, n),$$

with equality for $i = n$. In other words, the Hodge and Newton polygons have the same endpoints, and the Newton polygon is everywhere on or above the Hodge polygon.

Proof. Again, the case $i = 1$ is clear because σ_1 is the operator norm of A, and the general case follows by considering exterior powers. □

Again, equality has a structural meaning, but the proof requires a bit more work than in the archimedean case, since we no longer have access to orthogonality.

Theorem 17 (Hodge-Newton decomposition)**.** *Suppose that for some $i \in \{1, \dots, n-1\}$ we have*

$$\sigma_i > \sigma_{i+1}, \qquad |\lambda_i| > |\lambda_{i+1}|,$$
$$\sigma_1 \cdots \sigma_i = |\lambda_1 \cdots \lambda_i|.$$

(That is, the Hodge and Newton polygons share a vertex with x-coordinate $-n+i$.) Then there exists $U \in \mathrm{GL}_n(\mathfrak{o}_F)$ such that $U^{-1}AU$ is block diagonal, with the first block accounting for the first i singular values and eigenvalues, and the second block accounting for the others.

Proof. Let $v_1, \dots, v_n$ be a basis of F^n such that $v_1, \dots, v_i$ span the generalized eigenspaces with eigenvalues $\lambda_1, \dots, \lambda_i$. and $v_{i+1}, \dots, v_n$ span the generalized eigenspaces with eigenvales $\lambda_{i+1}, \dots, \lambda_n$. Apply the Smith normal form to construct a basis $w_1, \dots, w_n$ of $\mathfrak{o}_K^n$ such that $|Aw_i| = \sigma_i |w_i|$.

Since $\sigma_i > \sigma_{i+1}$, the only vectors $v \in \wedge^i \mathbb{C}^n$ for which $|Av|/|v|$ achieves its maximum value $\sigma_1 \cdots \sigma_i$ are the nonzero multiples of $w_1 \wedge \cdots \wedge w_i$. However, this is also true for $v_1 \wedge \cdots \wedge v_i$. We conclude that $w_1, \dots, w_i$ span V; this implies that we can conjugate A by a matrix in $\mathrm{GL}_n(\mathfrak{o}_F)$ into block diagonal form

$$\begin{pmatrix} B & C \\ 0 & D \end{pmatrix},$$

where B accounts for the first i Hodge and Newton slopes of A.

Each $i \times i$ minor of the matrix (BC) has valuation at least the sum of the first i Hodge slopes of A, which is the valuation of $\det(B)$. By Cramer's rule, each column of C is a $\mathfrak{o}_F$-linear combination of the columns of B, i.e., $B^{-1}C$ has entries in $\mathfrak{o}_F$. Moreover, the first Hodge slope of D is greater than the last Hodge slope of B, so $|B^{-1}CD| < |C|$. Thus conjugating by the matrix

$$\begin{pmatrix} I_i & -B^{-1}C \\ 0 & 1 \end{pmatrix}$$

gives a new matrix

$$\begin{pmatrix} B & C_1 \\ 0 & D \end{pmatrix}$$

with $|C_1| < C$. Repeating, we converge to a change of basis which converts A into the block diagonal matrix

$$\begin{pmatrix} B & 0 \\ 0 & D \end{pmatrix}$$

which has the desired form. $\square$

Note that the slopes of the Hodge polygon are forced to be in the (additive) value group of F, whereas the slopes of the Newton polygon need only lie in the divisible closure of the value group. Consequently, it is possible for a matrix to have no conjugates over $\mathrm{GL}_n(F)$

for which the Hodge and Newton polygons coincide. However, the following is true; see also Proposition 24 below.

Lemma 18. *Suppose that one of the following holds.*

(a) The value group of $|F^*|$ is dense in $\mathbb{R}_{>0}$, and $\eta > 1$.

(b) We have $\sigma_i \in |F^*|$ for $i = 1, \dots, n$ (so in particular $\sigma_i > 0$), and $\eta \geq 1$.

Then there exists $U \in \mathrm{GL}_n(F)$ such that the i-th singular value of $U^{-1}AU$ is at most $\eta|\lambda_i|$ (with equality in case (b)).

Proof. Case (b) is directly analogous of Lemma 8. We will prove a stronger form of (a) in a later unit. □

4 Perturbations (nonarchimedean case)

Again, we can ask about the effect of perturbations. The analogue of Weyl's second inequality is more or less trivial.

Proposition 19. *If B is a matrix with $|B| < \sigma_i$, then the first i singular values of $A + B$ are $\sigma_1, \dots, \sigma_i$.*

Proof. Exercise. □

We next consider the effect on the characteristic polynomial.

Theorem 20. *Let B be an $n \times n$ matrix such that $|B| < \sigma_j$ for some $j \in \{1, \dots, n\}$. Let $P(T) = T^n + \sum_{i=0}^{n-1} P_i T^i$ and $Q(T) = T^n + \sum_{i=0}^{n-1} Q_i T^i$ be the characteristic polynomials of A and $A + B$. Then*

$$|P_{n-i} - Q_{n-i}| \leq \sigma_1 \cdots \sigma_{i-1} |B| \qquad (i = 1, \dots, j).$$

Proof. The proof is as for Theorem 10, except now the factor $|2^n \binom{n}{i}^2|$ is dominated by 1. □

Question 21. *Is Theorem 20 best possible?*

We may also consider multiplicative perturbations.

Proposition 22. *Let $B \in \mathrm{GL}_n(F)$ satisfy $|B| \leq \eta$. Let $\sigma'_1, \dots, \sigma'_n$ be the singular values of AB. Then*

$$\sigma'_i \leq \eta \sigma_i \qquad (i = 1, \dots, n).$$

Proof. As for Proposition 11, but using the Smith normal form instead of the singular value decomposition. □

Corollary 23. *Suppose that the Newton and Hodge slopes of A coincide, and that $U \in \mathrm{GL}_n(F)$ satisfies $|U| \cdot |U^{-1}| \leq \eta$. Then each Newton slope of $U^{-1}AU$ is at most $\log \eta$ more than the corresponding Hodge slope.*

Here is a weak converse to Corollary 23. (We leave the archimedean analogue to the reader's imagination.)

Proposition 24. *Suppose that the Newton slopes of A are nonnegative and that $\sigma_1 \geq 1$. Then there exists $U \in \mathrm{GL}_n(F)$ such that*

$$|U^{-1}AU| \leq 1, \quad |U^{-1}| \leq 1, \quad |U| \leq \sigma_1^{n-1}.$$

Proof. Let $e_1, \dots, e_n$ denote the standard basis vectors. Let M be the smallest $\mathfrak{o}_F$-submodule of F^n containing $e_1, \dots, e_n$ and stable under A. For each i, if $j = j(i)$ is the least integer such that $e_i, Ae_i, \dots, A^j e_i$ are linearly dependent, then we have $A^j e_i = \sum_{h=0}^{j-1} c_h A^h e_i$ for some $c_h \in F$, and the nonnegativity of the Newton slopes forces $|c_h| \leq 1$. Hence M is finitely generated, and thus free, over $\mathfrak{o}_F$.

Let $v_1, \dots, v_n$ be a basis of M, and let U be the change-of-basis matrix $v_j = \sum_i U_{ij} e_i$; then $|U^{-1}AU| \leq 1$ because M is stable under A, and $|U^{-1}| \leq 1$ because M contains $e_1, \dots, e_n$. The desired bound on U will follow from the fact that for any $x = c_1 e_1 + \cdots + c_n e_n \in M$, we have

$$\max_i \{|c_i|\} \leq \sigma_1^{n-1}. \tag{1}$$

It suffices to check (1) for $x = A^h e_i$ for $i = 1, \dots, n$ and $h = 0, \dots, j(i) - 1$, as these generate M over $\mathfrak{o}_F$. But it is evident that $|A^h e_1| \leq \sigma_1^h |e_1| = \sigma_1^h$, so we are done. □

In one sense, examples like

$$A = \begin{pmatrix} 1 & c & 0 \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix}$$

with $|c| > 1$ show that this bound is sharp. However, one can get a slight improvement in special cases by accounting for the other singular values; see exercises.

By imitating the proof of Proposition 12, we obtain the following.

Proposition 25. *Let $\sigma_{k,1}, \dots, \sigma_{k,n}$ be the singular values of A^k. Then*

$$\lim_{k \to \infty} \sigma_{k,i}^{1/k} = |\lambda_i| \qquad (i = 1, \dots, n).$$

5 Horn's inequalities

Although they will not be needed in this course, it is quite natural to mention here some stronger versions of the perturbation inequalities in the archimedean and nonarchimedean cases, introduced by Horn [Hor62] in the archimedean case. See the beautiful survey article of Fulton [Ful00] for more information.

To introduce the stronger inequalities, we must set up some notation. Put

$$U_r^n = \{(I, J, K) : I, J, K \subseteq \{1, \dots, n\}, \#I = \#J = \#K = r,$$
$$\sum_{i \in I} i + \sum_{j \in J} j = \sum_{k \in K} k + r(r+1)/2\}.$$

For $(I, J, K) \in U_r^n$, write $I = \{i_1 < \cdots < i_r\}$ and similarly for J, K. For $r = 1$, put $T_1^n = U_1^n$. For $r > 1$, put

$$T_r^n = \{(I, J, K) \in U_r^n : \text{for all } p < r \text{ and } (F, G, H) \in T_p^r,$$
$$\sum_{f \in F} i_f + \sum_{g \in G} j_g \leq \sum_{h \in H} k_h + p(p+1)/2\}.$$

For multiplicative perturbations, we obtain the following results. which include the Weyl inequalities (Theorem 4, Theorem 16) as well as Propositions 11 and 22. (It is important for the proofs that one can rephrase the Horn inequalities in terms of Littlewood-Richardson numbers; see [Ful00, §3].)

Theorem 26. *For $* \in \{A, B, C\}$, let $\sigma_{*,1}, \ldots, \sigma_{*,n}$ be a nonincreasing sequence of nonnegative real numbers. Then the following are equivalent.*

(a) There exist $n \times n$ matrices A, B, C over $\mathbb{C}$ with $AB = C$ such that for $* \in \{A, B, C\}$, $*$ has singular values $\sigma_{*,1}, \ldots, \sigma_{*,n}$.

(b) We have $\prod_{i=1}^n \sigma_{A,i} \prod_{j=1}^n \sigma_{B,j} = \prod_{k=1}^n \sigma_{C,k}$, and for all $r < n$ and $(I, J, K) \in T_r^n$, we have
$$\prod_{k \in K} \sigma_{C,k} \leq \prod_{i \in I} \sigma_{A,i} \prod_{j \in J} \sigma_{B,j}.$$

Proof. See [Ful00, Theorem 16]. Note that the first condition in (b) is omitted in the statement given in [Ful00], but this is only a typo. □

Theorem 27. *Let F be a complete nonarchimedean field with additive value group G. For $* \in \{A, B, C\}$, let $\sigma_{*,1}, \ldots, \sigma_{*,n}$ be a nonincreasing sequence of elements of $G \cup \{0\}$. Then the following are equivalent.*

(a) There exist $n \times n$ matrices A, B, C over F with $AB = C$ such that for $* \in \{A, B, C\}$, $*$ has singular values $\sigma_{*,1}, \ldots, \sigma_{*,n}$.

(b) We have $\prod_{i=1}^n \sigma_{A,i} \prod_{j=1}^n \sigma_{B,j} = \prod_{k=1}^n \sigma_{C,k}$, and for all $r < n$ and $(I, J, K) \in T_r^n$, we have
$$\prod_{k \in K} \sigma_{C,k} \leq \prod_{i \in I} \sigma_{A,i} \prod_{j \in J} \sigma_{B,j}.$$

Proof. See [Ful00, Theorem 7]. □

For additive perturbations, one has an analogous result in the archimedean case; see [Ful00, Theorem 15]. I am not aware of an additive result in the nonarchimedean case. Also, in the archimedean case one has analogous results (with slightly different statements) in which one restricts to Hermitian matrices.

6 Notes

See [Bha97, §III] for results in the archimedean case not otherwise cited, such as the fact that a real symmetric matrix has nonnegative real eigenvalues, and the singular value decomposition. (This book was out of the library when I wrote this, so I wasn't able to look up precise references.) We unfortunately cannot recommend a good reference for the strong analogy between archimedean and p-adic numerical analysis; this seems to be a poorly known piece of folklore.

In Theorem 7, the equivalence of (a) and (b) is standard. We do not have a reference for the equivalence with (c), although it is implicit in most proofs of the equivalence of (a) and (b).

The reader familiar with the notions of elementary divisors or invariant factors may be wondering why the terminology "Hodge polygon" is necessary or reasonable. The answer is that the Hodge numbers of a variety over a p-adic field are reflected by the elementary divisors of the action of Frobenius on crystalline cohomology. The fact that the Newton polygon lies above the Hodge polygon then implies a relation between the characteristic polynomial of Frobenius and the Hodge numbers of the original variety; this relationship was originally conjectured by Katz and proved by Mazur. See [BO78] for further discussion of this point, and of crystalline cohomology as a whole.

Much of the work in this chapter can be carried over to the case of a transformation which is only semilinear for some isometric endomorphism of F. This case arises in the study of slope filtrations of Frobenius crystals (F-crystals), as in [Kat79]; in fact, the Hodge-Newton decomposition theorem (Theorem 17) is a direct translation of Katz's corresponding theorem for F-crystals [Kat79, Theorem 1.6.1]. The archimedean version (Theorem 6) is itself a translation of Theorem 17; we do not know of a reference, although we do not make any claim of originality. Likewise, Proposition 25 is a direct translation of [Kat79, Corollary 1.4.4]; its archimedean analogue (Proposition 12) is doubtless also known, but we do not have a reference.

The question of how much the characteristic polynomial of a square matrix over a field is affected by a perturbation arises in numerical applications. This is a familiar fact in the archimedean case, but perhaps less so in the nonarchimedean case; numerical applications of the latter include using p-adic cohomology to compute zeta functions of varieties over finite fields. See for instance [AKR07, §1.6], [Ger07, §3].

7 Exercises

1. Prove Theorem 13. (Hint: first do row reduction to get an upper triangular matrix whose diagonal entries have norms equal to the singular values.)

2. Prove the following nonarchimedean analogue of Horn's theorem (Theorem 5): any $\lambda_1, \dots, \lambda_n \in F$ and $\sigma_1, \dots, \sigma_n \in 0 \cup G$ (where G is the additive value group of F) satisfying Weyl's inequalities (including the equality for $i = n$) occur as the eigenvalues and singular values of an $n \times n$ matrix over F.

3. Prove Proposition 19.
4. Let $D \in \mathrm{GL}_n(F)$ be a diagonal matrix, and let $U, V \in \mathrm{GL}_n(\mathfrak{o}_F)$ be congruent to the identity matrix modulo $\mathfrak{m}_F$. Prove that the Newton polygons of D and UDV coincide. This is [BC05, Lemma 5]. Optional: is there an archimedean analogue?
5. Improve the upper bound in (1), and consequently in Proposition 24, to $\prod_{i=1}^{n-1} \max\{1, \sigma_i\}$. (Hint: reduce to the case where A admits a cyclic vector.) Optional: is there an archimedean analogue?