NOTES ON THE PERRON-FROBENIUS THEORY OF NONNEGATIVE MATRICES

MIKE BOYLE

1. Introduction

By a nonnegative matrix we mean a matrix whose entries are nonnegative real numbers. By positive matrix we mean a matrix all of whose entries are strictly positive real numbers.

These notes give the core elements of the Perron-Frobenius theory of nonnegative matrices. This splits into three parts:

(1) the primitive case (due to Perron)
(2) the irreducible case (due to Frobenius)
(3) the general case (due to?)

We will skip the (in most respects easy) transition from the irreducible to the general case.

2. The primitive case

Definition 2.1. A *primitive* matrix is a square nonnegative matrix some power of which is positive.

The primitive case is the heart of the Perron-Frobenius theory and its applications. There are various proofs. See the final remarks for acknowledgments on this one.

The spectral radius of a square matrix is the maximum of the moduli of its eigenvalues. A number λ is a *simple* root of a polynomial $p(x)$ if it is a root of multiplicity one (i.e., $p(\lambda) = 0$ and $p'(\lambda) \neq 0$). For a matrix A or vector v, we define the norm ($||A||$ or $||v||$) to be the sum of the absolute values of its entries.

Theorem 2.2 (Perron Theorem)**.** *Suppose A is a primitive matrix, with spectral radius λ. Then λ is a simple root of the characteristic polynomial which is strictly greater than the modulus of any other root, and λ has strictly positive eigenvectors.*

For example, the matrix $\begin{pmatrix} 0 & 2 \\ 1 & 1 \end{pmatrix}$ is primitive (with eigenvalues $2, -1$), but the matrices $\begin{pmatrix} 0 & 4 \\ 1 & 0 \end{pmatrix}$ (with eigenvalues $2, -2$) and $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ (with 1 a repeated eigenvalue) are not. Note that the "simple root" condition is stronger than the condition that λ have a one dimensional eigenspace, because a one-dimensional eigenspace may be part of a larger-dimensional generalized eigenspace.

We begin with a geometrically compelling lemma.

Lemma 2.3. *Suppose T is a linear transformation of a finite dimensional vector space, S' is a polyhedron containing the origin in its interior, and a positive power of T maps S' into its interior. Then the spectral radius of T is less than 1.*

Proof of the lemma. Without loss of generality, we may suppose T maps S' into its interior. Clearly, there is no eigenvalue of modulus greater than 1.

The image of S' is a closed set which does not intersect the boundary of S'. Because $T^n(S') \subset T(S')$ if $n \geq 1$, no point on the boundary of S' can be an image of a power of T, or an accumulation point of points which are images of powers of T. But this is contradicted if T has an eigenvalue of modulus 1, as follows:

CASE I: a root of unity is an eigenvalue of a T.

In this case, 1 is an eigenvalue of a power of T, and a power of T has a fixed point on the boundary of S'. Thus the image of S' under a power of T intersects the boundary of S', a contradiction.

CASE II: there is an eigenvalue of modulus 1 which is not a root of unity.

In this case, let V be a 2-dimensional subspace on which T acts as an irrational rotation. Let p be a point on the boundary of S' which is in V. Then p is a limit point of $\{T^n(p) : n > 1\}$, so p is in the image of T, a contradiction.

This completes the proof. □

Proof of the Perron Theorem. There are three steps.

STEP 1: get the positive eigenvector.

The unit simplex S is the set of nonnegative vectors v such that $||v|| = 1$. The matrix A induces the continuous map from S into itself which sends a vector v to $||v||^{-1}v$. By Brouwer's Fixed Point Theorem, this map has a fixed point, which shows there exists a nonnegative eigenvector. Because a power of A is positive, the eigenvector must actually be positive. Let λ be the eigenvalue, which is positive.

STEP 2: stochasticize A.

Let r be a positive right eigenvector. Let R be the diagonal matrix whose diagonal entries come from r, i.e. $R(i,i) = r_i$. Define the matrix $P = (1/\lambda)R^{-1}AR$. P is still primitive. The column vector with every entry equal to 1 is an eigenvector of P with eigenvalue 1, i.e. P is stochastic. It now suffices to do Step 3.

STEP 3: show 1 is a simple root of the characteristic polynomial of P dominating the modulus of any other root.

Consider the action of P on row vectors: P maps the unit simplex S into itself and a power of P maps S into its interior. From Step 1, we know there is a positive row vector v in S which is fixed by P. Therefore $S' = -v + S$ is a polyhedron, whose interior contains the origin. By the lemma the restriction of P to the subspace V spanned by S' has spectral radius less than 1. But V is P-invariant with codimension 1. □

We can now check that a primitive matrix has (up to scalar multiples) just one nonnegative eigenvector.

Corollary 2.4. *Suppose A is a primitive matrix and w is a nonnegative vector, with eigenvalue β. Then β must be the spectral radius of A.*

Proof. Because A is primitive, we can choose $k > 0$ such that $A^k w$ is positive. Thus, $w > 0$ (since $A^k w = \beta^k w$) and $\beta > 0$. Now choose a positive eigenvector v which has eigenvalue λ, the spectral radius, such that $v < w$. Then for all $n > 0$,

$$\lambda^n v = A^n v \leq A^n w = \beta^n w \ .$$

This is impossible if $\beta < \lambda$, so $\beta = \lambda$. □

Remarks 2.5. (1) Any number of people have noticed that applicability of Brouwer's Theorem. (Charles Johnson told me Ky Fan did this in the 1950's.) It's a matter of taste as to whether to use it to get the eigenvector.

(2) The proof above, using the easy reduction to the geometrically clear and simple lemma, was found by Michael Brin in 1993. It is dangerous in this area to claim a proof is new, and I haven't read the German papers of Perron and Frobenius themselves. However I haven't seen this reduction in the other proofs of the Perron theorem I've read.

(3) The utility of the stochasticization trick is by no means confined to this theorem.

3. The irreducible case

Given a nonnegative $n \times n$ matrix A, we let its rows and columns be indexed in the usual way by $\{1, 2, \dots n\}$, and we define a directed graph $G(A)$ with vertex set $\{1, 2, \dots, n\}$ by declaring that there is an edge from i to j if and only if $A(i,j) \neq 0$. A loop of length k in $G(A)$ is a path of length k (a path of k successive edges) which begins and ends at the same vertex.

Definition 3.1. An irreducible matrix is a square nonnegative matrix such that for every i, j there exists $k > 0$ such that $A^k(i,j) > 0$.

Notice, for any positive integer k, $A^k(i,j) > 0$ if and only if there is a path of length k in $G(A)$ from i to j.

Definition 3.2. The *period* of an irreducible matrix A is the greatest common divisor of the lengths of loops in $G(A)$.

For example, the matrix $\begin{pmatrix} 0 & 2 \\ 1 & 1 \end{pmatrix}$ has period 1 and the matrix $\begin{pmatrix} 0 & 4 \\ 1 & 0 \end{pmatrix}$ has period 2.

Now suppose A is irreducible with period p. Pick some vertex v, and for $0 \leq i, p$ define a set of vertices

$$C_i = \{u : \text{there is a path of length } n \text{ from } v \text{ to } u \text{ such that } n \equiv i \mod p\} .$$

The sets $C(i)$ partition the vertex set. An arc from a vertex in $C(i)$ must lead to a vertex in $C(j)$ where $j = i + 1 \mod p$. If we reorder the indices for rows and columns of A, listing indices for C_0, then C_l, etc., and replace A with PAP^{-1} where P is the corresponding permutation matrix, then we get a matrix B with a block form which looks like a cyclic permutation matrix. For example, with $p = 4$, we have a block matrix

$$B = \begin{pmatrix} 0 & A_1 & 0 & 0 \\ 0 & 0 & A_2 & 0 \\ 0 & 0 & 0 & A_3 \\ A_4 & 0 & 0 & 0 \end{pmatrix} .$$

An specific example with $p = 3$ is

$$\begin{pmatrix} 0 & 2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 1 & 2 \\ 0 & 0 & 0 & 1 & 0 & 3 \\ 3 & 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} .$$

Note the blocks of B are rectangular (not necessarily square). B and A agree on virtually all interesting properties, so we usually just assume A has the form given as B (i.e., we tacitly replace A with B, not bothering to rename). We call this a cyclic block form.

Proposition 3.3. *Let A be a square nonnegative matrix. Then A is primitive if and only if it is irreducible with period one.*

Proof. Exercise. □

Definition 3.4. We say two matrices have the same nonzero spectrum if their characteristic polynomials have the same nonzero roots, with the same multiplicities.

Proposition 3.5. *Let A be an irreducible matrix of period p in cyclic block form. Then A^p is a block diagonal matrix and each of its diagonal blocks is primitive. Moreover each diagonal block has the same nonzero spectrum.*

Proof. These diagonal blocks must be irreducible of period 1, hence primitive. Each has the form $D(i) = A(i)A(i+l)\cdots A(i+p)$ where $A(j)$ is the nonzero block in the jth block row and j is understood mod p. Thus given i there are rectangular matrices S, R such that $D(i) = SR$, $D(i+l) = RS$. Therefore their nth powers are $S((RS)^{n-1}R)$ and $((RS)^{n-1}R)S$, so for each n their nth powers have the same trace (because $\mathrm{trace}(UV) = \mathrm{trace}(VU)$ for any matrices U, V). This forces $D(i)$ and $D(i+1)$ to have the same nonzero spectrum. (In fact the nonnilpotent part of the Jordan form for $D(i)$ is the same for all i.) □

Proposition 3.6. *Let A be an irreducible matrix with period p and suppose that ξ is a primitive pth root of unity. Then the matrices A and ξA are similar.*

In particular, if c is root of the characteristic polynomial of A with multiplicity m, then ξc is also a root with multiplicity m.

Proof. The proof for the period 3 case already explains the general case:

$$\begin{pmatrix} \xi^{-1}I & 0 & 0 \\ 0 & \xi^{-2}I & 0 \\ 0 & 0 & \xi^{-3}I \end{pmatrix} \begin{pmatrix} 0 & A_1 & 0 \\ 0 & 0 & A_2 \\ A_3 & 0 & 0 \end{pmatrix} \begin{pmatrix} \xi^{1}I & 0 & 0 \\ 0 & \xi^{2}I & 0 \\ 0 & 0 & \xi^{3}I \end{pmatrix}$$

$$= \begin{pmatrix} 0 & \xi A_1 & 0 \\ 0 & 0 & \xi A_2 \\ \xi^{-2}A_3 & 0 & 0 \end{pmatrix} = \xi \begin{pmatrix} 0 & A_1 & 0 \\ 0 & 0 & A_2 \\ A_3 & 0 & 0 \end{pmatrix}$$

since $\xi^{-2} = \xi$. □

Definition 3.7. If A is a matrix, then its characteristic polynomial away from zero is the polynomial $q_A(t)$ such that $q_A(0)$ is not 0 and the characteristic polynomial of A is a power of t times $q_A(t)$.

Theorem 3.8. *Let A be an irreducible matrix of period p. Let D be a diagonal block of A^p (so, D is primitive). Then*

$$q_A(t) = q_D(t^p) \ .$$

Equivalently, if ξ is a primitive pth root of unity and we choose complex numbers $\lambda_1, \ldots, \lambda_j$ such that $q_D(t) = \prod_{j=1}^{k}(t - (\lambda_j^p))$, then

$$q_A(t) = \prod_{i=0}^{p-1}\prod_{j=1}^{k}(t - \xi^i \lambda_j) \ .$$

Proof. From the last proposition, a nonzero root c of q_{A^p} has multiplicity kp, where k is the number such that every pth root of c is a root of multiplicity k of q_A. Each c which is a root of multiplicity k for q_D is a root of multiplicity kp for q_{A^p} (since the diagonal blocks of A^p have the same nonzero spectrum. $\square$

Theorem 3.9 (Perron-Frobenius Theorem)**.** *Let A be an irreducible matrix of period p.*

(1) *A has a nonnegative right eigenvector r. This eigenvector is strictly positive, its eigenvalue λ is the spectral radius of A, and any nonnegative eigenvector of A is a scalar multiple of r.*
(2) *The roots of the characteristic polynomial of A of modulus λ are all simple roots, and these roots are precisely the p numbers $\lambda, \xi\lambda, \ldots, \xi^{p-1}\lambda$ where ξ is a primitive pth root of unity.*
(3) *The nonzero spectrum of A is invariant under multiplication by ξ.*

Proof. Everything is easy from what has gone before except the construction of the eigenvector. The general idea is already clear for $p = 3$. Then we can consider A in the block form

$$A = \begin{pmatrix} 0 & A_1 & 0 \\ 0 & 0 & A_2 \\ A_3 & 0 & 0 \end{pmatrix} .$$

Now $A_1A_2A_3$ is a diagonal block of A, primitive with spectral radius λ^3. Let r be a positive right eigenvector for $A_1A_2A_3$. Compute:

$$\begin{pmatrix} 0 & A_1 & 0 \\ 0 & 0 & A_2 \\ A_3 & 0 & 0 \end{pmatrix} \begin{pmatrix} \lambda^2 r \\ A_2A_3r \\ \lambda A_3 r \end{pmatrix} = \begin{pmatrix} A_1A_2A_3r \\ \lambda A_2A_3r \\ \lambda^2 A_3 r \end{pmatrix} = \lambda \begin{pmatrix} \lambda^2 r \\ A_2A_3r \\ \lambda A_3 r \end{pmatrix}$$

$\square$

4. An application

The Perron theorem provides a very clear picture of the way large powers of a primitive matrix behave.

Theorem 4.1. *Suppose A is primitive. Let u be a positive left eigenvector and let v be a positive right eigenvector for the spectral radius λ, chosen such that $uv = 1$. Then $((1/\lambda)A)^n$ converges to the matrix vu, exponentially fast.*

Remark 4.2. The theorem says that for large n, $A^n - \lambda^n vu$ has entries much smaller than A^n; the dominant behavior of A^n is described by the simple matrix $\lambda^n vu$. For example, if $A = \begin{pmatrix} 1 & 3 \\ 2 & 2 \end{pmatrix}$, then A has spectral radius $\lambda = 4$, with left and right eigenvectors $(2, 3)$ and $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$. Their inner product is 5; so, we can take

$$u = (2/5, 3/5) , \quad v = \begin{pmatrix} 1 \\ 1 \end{pmatrix} , \quad \text{and} \quad vu = \begin{pmatrix} 2/5 & 3/5 \\ 2/5 & 3/5 \end{pmatrix} .$$

One can check that indeed

$$A^n = 4^n \begin{pmatrix} 2/5 & 3/5 \\ 2/5 & 3/5 \end{pmatrix} + (-1)^n \begin{pmatrix} 3/5 & -3/5 \\ -2/5 & 2/5 \end{pmatrix} .$$

Proof of theorem. The matrix $(1/\lambda)A$ multiplies the eigenvectors u and v by 1 (i.e. leaves them unchanged). Now suppose w is a generalized column eigenvector for A for eigenvalue β. By the Perron Theorem we have $|\beta| < \lambda$, so $\lim_n((1/\lambda)A)^n w$ converges to the zero vector (exponentially fast). The same holds for row vectors. Consequently $((1/\lambda)A)^n$ converges to a matrix M which fixes u and v and which annihilates the other generalized eigenvectors of A. This matrix is unique. We claim that $M = vu$. For this we first note that $uM = u(vu) = (uv)u = u$ and similarly $Mv = v$. Now suppose w is a generalized column eigenvector for eigenvalue β not equal to λ: then $Mw = 0$, because otherwise $Mw = vuw \neq 0$ would imply $uw \neq 0$ and then for all $n > 0$

$$\lambda^n uw = (uA^n)w = u(A^n w)$$

so

$$uw = (1/\lambda^n)u(A^n w)$$

which is impossible because

$$\lim_{n\to\infty} \frac{1}{\lambda^n} A^n w = 0 \ .$$

Similarly, $wM = 0$ if w is a generalized row eigenvector for A for eigenvalue other than λ.

$\square$

Department of Mathematics, University of Maryland, College Park, MD 20742-4015, U.S.A.

E-mail address: mmb@math.umd.edu

URL: www.math.umd.edu/∼mmb