PERRON-FROBENIUS THEOREM

Definition 0.1. A matrix or vector is non-negative (positive) if all of its entries are non-negative (positive). A non-negative square matrix A such that $A^k > 0$ for some $k \in \mathbb{N}$ is called a primitive matrix (also known as regular matrix). The directed graph D_A corresponding to a non-negative square matrix A is the digraph whose vertices are the coordinates, and (i, j) is a directed edge iff $A[i, j] > 0$. We say that the square matrix $A \geq 0$ is irreducible if the corresponding digraph D_A is strongly connected.

Recall a standard fact from linear algebra (or in fact more generally from functional analysis).

Proposition 0.2 (Gelfand)**.** *Let ρ be the spectral radius of the square matrix A. Then*

$$\rho = \lim_{m\to\infty} ||A^m||^{1/m}$$

We need the following elementary lemma.

Lemma 0.3. *Let $A > 0$ be a square matrix, and $\underline{u}, \underline{v}$ with $\underline{u} - \underline{v} \geq 0$ but not equal to the zero vector. Then for some $C > 1$ we have $A\underline{u} > CA\underline{v}$.*

Proof. Clearly $A(\underline{u} - \underline{v}) > 0$, thus $A\underline{u} > A\underline{v}$, and the assertion follows trivially. $\square$

Theorem 0.4 (Perron-Frobenius, primitive case)**.** *Let A be an $n \times n$ primitive matrix with spectral radius ρ. Then*

(1) *ρ is an eigenvalue with associated eigenvector $\underline{v} > 0$ (Perron root, Perron vector);*
(2) *for all other eigenvalues $\lambda \in \mathrm{Spec}(A)$ we have $|\lambda| < \rho$;*
(3) *ρ has algebraic multiplicity 1.*
(4) *If $\underline{w}$ is the Perron vector for A^*, that is, $\underline{w}^*$ is the left Perron vector for A, such that $\underline{w}^*$ and $\underline{v}$ are normalized to satisfy the identity $\underline{w}^*\underline{v} = 1$, then $\lim_{m\to\infty} \left(\frac{1}{\rho}A\right)^m = \underline{vw}^*$.*

Proof. We prove the first three items for **positive** matrices first: so let us assume that $A > 0$. Let $\lambda \in \mathrm{Spec}(A)$ be an eigenvalue with $|\lambda| = \rho$ and eigenvector $\underline{u}$. Let $\underline{v} = |\underline{u}|$ be the vector obtained by taking the absolute value of $\underline{u}$ coordinate-wise. The triangle inequality implies $A\underline{v} \geq |A\underline{u}| = |\lambda\underline{u}| = \rho\underline{v}$. Assume indirectly that equality does not hold. Then by Lemma 0.3 there is a $C > 1$ such that $A^2\underline{v} > (C\rho)A\underline{v}$. A standard induction yields $A^m(A\underline{v}) = A^{m+1}\underline{v} > (C\rho)^m A\underline{v}$. By Gelfand's formula we obtain $\rho \geq C\rho$, a contradiction.

Hence, $A\underline{v} = \rho\underline{v}$. To finish the proof of item (1), we need to show that $\underline{v} > 0$: by definition, it is non-negative and not the zero vector, but in principle, it could have some zero entries. However, the left-hand side $A\underline{v}$ is positive, so the right-hand side $\rho\underline{v}$ must be positive, too.

For item (2), assume that there is a $\rho \neq \lambda \in \mathrm{Spec}(A)$ with $|\lambda| = \rho$; let $\underline{u}$ be an associated eigenvector. The above calculation shows that $A|\underline{u}| = |A\underline{u}|$. Thus in every coordinate, the absolute value of a positive linear combination of complex numbers equals to the linear combination of absolute values. It is an easy exercise to show that this is only possible if the complex numbers are all non-negative multiples of the same number $z \in \mathbb{C}$. Hence, $\underline{u} = z\underline{u}'$ for some $\underline{u}' \geq 0$. Then we may assume that $\underline{u} = \underline{u}'$, as $\underline{u}'$ is also an eigenvector with eigenvalue λ. As $A\underline{u} = \lambda\underline{u}$, and the left-hand side is positive, we obtain $\lambda > 0$. As $|\lambda| = \rho$ and $\lambda > 0$, we have $\lambda = \rho$.

In order to show item (3), we first prove the formally weaker statement that the geometric multiplicity of ρ is 1. To this end, assume that there is an eigenvector $\underline{u}$ associated to ρ such that $\underline{u}$ and $\underline{v}$ are linearly independent. The coordinate-wise real and imaginary part of $\underline{u}$ are also eigenvectors associated to the real eigenvalue ρ. If both were a multiple of $\underline{v}$, then $\underline{u}$ and $\underline{v}$ would not be independent. Thus we can switch $\underline{u}$ to one of these real vectors, and assume that $\underline{u} \in \mathbb{R}^n$. We may also assume that at least one entry of $\underline{u}$ is positive. Then a standard continuity argument yields a $c > 0$ such that $\underline{v} - c\underline{u} \geq 0$ with at least one zero entry. By the independence of $\underline{u}$ and $\underline{v}$, the vector $\underline{v} - c\underline{u} \neq \underline{0}$. However, $\rho(\underline{v} - c\underline{u}) = A(\underline{v} - c\underline{u}) > 0$, a contradiction.

Let $\underline{w}^*$ be the left Perron vector for A. Let $U \leq \mathbb{R}^n$ be the perpendicular subspace to $\underline{w}$, that is, $U = \{\underline{x} \in \mathbb{R}^n \mid \underline{w}^*\underline{x} = 0\}$. As $\underline{w}, \underline{v} > 0$, we have $\underline{v} \notin U$. Moreover, U is an $(n-1)$-dimensional A-invariant subspace: if $\underline{x} \in U$, then $\underline{w}^*(A\underline{x}) = (\underline{w}^*A)\underline{x} = \rho\underline{w}^*\underline{x} = 0$.

Let us pick a new basis of $\mathbb{R}^n$: the first basis vector is $\underline{v}$, and the remaining $(n-1)$ forms a basis of U. The transition matrix is denoted by S. Then the matrix after the base transition is $S^{-1}AS$. As $\langle\underline{v}\rangle$ and U are A-invariant subspaces, we have that $S^{-1}\langle\underline{v}\rangle$ and $S^{-1}U$ are $S^{-1}AS$-invariant subspaces. So over the new basis, the matrix $S^{-1}AS$ is block diagonal, with a one-by-one block containing ρ. Hence, if the algebraic multiplicity of ρ is greater than 1, then an eigenvector in U would be associated to ρ, making its geometric multiplicity greater than 1.

Now that items $(1,2,3)$ are proven for positive matrices, we reduce the general statements for primitive matrices to this special case. Assuming that $A^k > 0$, let ρ^k be the Perron root (spectral radius) of A^k, and let $\underline{v}$ be the associated positive eigenvector. All other eigenvalues have strictly smaller modulus, and ρ^k has algebraic multiplicity 1. Then the eigenvalues of A are k-th roots of those of A^k with the same associated eigenvectors. Hence, $\rho\varepsilon$ is the spectral radius of A for some k-th root of unity ε, and it is also an eigenvalue with associated positive eigenvector $\underline{v}$, and $\rho\varepsilon$ has algebraic multiplicity 1. As $A\underline{v} = \rho\varepsilon\underline{v}$, and the left-hand side is positive (real), we conclude that $\varepsilon = 1$.

For item (4), let T be the transition matrix such that $T^{-1}(\frac{1}{\rho}A)T = J$ is the Jordan normal form of $\frac{1}{\rho}A$. Then $(\frac{1}{\rho}A)^m = (TJT^{-1})^m = TJ^mT^{-1} \to TMT^{-1}$, where M is the square matrix whose upper left element is 1 and all other elements are 0. The first column of T is a scalar multiple of $\underline{v}$. The first row of T^{-1} is a scalar multiple of $\underline{w}^*$: indeed, the same argument can be applied to A^*, and then the transition matrix is $(T^{-1})^*$. Thus we may assume that these scalar multiplies are $\underline{v}$ and $\underline{w}^*$ themselves, provided that $\underline{w}^*\underline{v} = 1$. Thus $TMT^{-1} = \underline{v}\underline{w}^*$. □

Remark 0.5. *According to the proof, $TMT^{-1} = \underline{v}\underline{w}^*$, thus $\underline{v}\underline{w}^*$ is a projection matrix with rank 1; that is, its left- and right images are 1-dimensional subspaces. As $\underline{w}^*\underline{v}\underline{w}^* = 1\underline{w}^* = \underline{w}^*$ and $\underline{v}\underline{w}^*\underline{v} = \underline{v}1 = \underline{v}$, the matrix $\underline{v}\underline{w}^*$ projects onto the 1-dimensional subspace(s) spanned by the Perron vector(s) (the plural being justified by the two actions of the matrix by left- and right multiplication). These are the Perron projections.*

Just as the primitive case was reduced to the positive one, we can reduce the irreducible version to the primitive one. Recall that the period of a vertex in a digraph is the greatest common divisor of all loops (walks with the same start- and endpoint) containing that vertex. In strongly connected digraphs, all vertices have the same period h, which we call the period of the digraph. This justifies the notion of a period of a non-negative irreducible matrix A: it is the period of the strongly connected digraph D_A. Note that such a digraph D_A (the set of coordinates of A) can be partitioned into h classes numbered from 0 to $h-1$, such that directed edges only go from class i to class $i+1$ modulo h. That is, A is a block diagonal matrix (after a permutation of coordinates), and nonzero elements can only be positioned in blocks with row class i and column class $i+1$ modulo h.

Theorem 0.6 (Perron-Frobenius, irreducible case)**.** *Let A be an $n \times n$ non-negative irreducible matrix with spectral radius ρ and period h. Then*

(1) *ρ is an eigenvalue with an associated eigenvector $\underline{v} > 0$ (Perron root, Perron vector), and for any h-th root of unity ε we have $\varepsilon\rho$ is also an eigenvalue;*
(2) *for all other eigenvalues $\lambda \in \operatorname{Spec}(A)$ we have $|\lambda| < \rho$;*
(3) *each $\varepsilon\rho$ has algebraic multiplicity 1.*
(4) *If $\underline{w}$ is the Perron vector for A^*, that is, $\underline{w}^*$ is the left Perron vector for A, such that $\underline{w}^*$ and $\underline{v}$ are normalized to satisfy the identity $\underline{w}^*\underline{v} = 1$, then $\lim_{m\to\infty} \frac{1}{m+1} \sum_{\ell=0}^{m} \left(\frac{1}{\rho}A\right)^{\ell} = \underline{vw}^*$.*

Proof. We use two constructions that make A a primitive matrix (or in fact, in the first case, a block diagonal matrix all of whose blocks are primitive). So let $P = A^h$; this is a block diagonal matrix. As D_A is strongly connected, and walks starting and ending in the same class have length divisible by h, we have that each block of P is primitive, as they correspond to the strongly connected components of D_A. Thus we can apply the previous theorem: all h blocks have a unique eigenvalue (with algebraic multiplicity 1) that has the largest absolute value (the spectral radius of the block). At this point, we do not know if the spectral radii of the blocks are equal; nevertheless, we already learned that there can be at most h eigenvalues of A (counted with multiplicity) whose absolute value is the spectral radius of A.

The other construction we consider is $Q = A + tI_n$, where I_n is the identity matrix and t is a small positive real number. Then the eigenvectors of A and Q coincide, and the associated eigenvalues are shifted by t. Clearly, D_Q contains all the directed edges of D_A, so D_Q is strongly connected. The common period of vertices in D_Q is 1, as we put a loop edge on every vertex by adding the identity matrix. Thus Q is an irreducible aperiodic matrix, hence primitive. In particular, the spectral radius of Q is a single eigenvalue $\rho + t$ with an associated positive eigenvector. Thus the slightly smaller value ρ is thus an eigenvalue of A with the same positive eigenvector $\underline{v}$. As this works for all $t > 0$, we have that ρ is the spectral radius of A, and its algebraic multiplicity is 1.

Let us partition the set of coordinates of $\underline{v}$ into the h blocks, creating the vectors $\underline{v}_0, \underline{v}_1, \dots, \underline{v}_{h-1}$. Let $\underline{v}'$ be the vector formed by the parts $\underline{v}_0, \varepsilon\underline{v}_1, \dots, \varepsilon^{h-1}\underline{v}_{h-1}$ for any h-th root of unity ε. Then $A\underline{v} = \rho\underline{v}$ and the sign pattern of blocks of A imply that $A_{i,i+1}\underline{v}_{i+1} = \rho\underline{v}_i$, where the indexation is to be understood modulo h, and $A_{i,i+1}$ is the block of elements in A whose row class is i and column class is $i+1$ (once again, modulo h). Thus $A_{i,i+1}\varepsilon^{i+1}\underline{v}_{i+1} = \rho\varepsilon^{i+1}\underline{v}_i = \varepsilon\rho(\varepsilon^i\underline{v}_i)$, showing that $A\underline{v}' = \varepsilon\rho\underline{v}'$. Hence, $\varepsilon\rho$ is an eigenvalue of A for every h-th root of unity ε. As we observed earlier that there can be at most h eigenvalues of A of length ρ, these must be all, finishing the proof of the first three items.

The fourth item can be shown similarly to item (4) in the primitive case. The difference is that now the Jordan normal form J of $\frac{1}{\rho}A$ has h one-by-one blocks, containing the numbers ε for every h-th root of unity ε; the remaining blocks still tend to 0 when taking large powers. The block containing 1 yields the constant 1 sequence when raised to different powers: so we have convergence there with limit 1. The problem occurs in the remaining blocks with entry having modulus 1. The powers of those complex h-th roots of unity form a divergent sequence. However, when averaged out, they tend to 0, as the sum of the first h powers of a complex number $\varepsilon \neq 1$ with $\varepsilon^h = 1$ is 0. $\square$