

642:550, Summer 2004, Supplement 7

A robust method for finding characteristic polynomial

Summer 2004

1. Introduction It is easy to read too much into the definition of the characteristic polynomial as a determinant. As soon as the n by n matrix has $n > 3$, it becomes extremely difficult to evaluate that determinant reliably. None of the methods for evaluating determinants work well when the matrix is allowed to contain an indeterminate. In this supplement, we describe an alternative method that can reliably obtain the characteristic polynomial of an integer matrix using hand computation for $n = 5$. It is also easy to program. A calculator version will be illustrated in lecture, and studies of the method from the middle of the twentieth century note that $n = 10$ was within reach on the digital computers of that time. However, if you are looking for eigenvalues and eigenvectors, an **approximate** method that concentrates on finding **some** eigenvectors to **high accuracy** is likely to be more useful in practice than this method, which can find eigenvalues only as roots of the characteristic polynomial.

The need for such a method was impressed on me when I wanted to compute a certain 5 by 5 integer determinant by hand and every method I tried gave me a different answer because of mistakes made while exercising the large number of arithmetic steps. Although the determinant is only one coefficient in the characteristic polynomial, it doesn't hurt to find more than is required. Since the calculations are less demanding and the method includes error-detection, this method is ideal for hand computation with integer matrices.

More details on the method can be found in F. R. Gantmacher, *The Theory of Matrices*, Chelsea, 1990, § IV.5. This follows the description by D. K. Faddeev given in the 1940s, although the method is attributed to work of Leverrier in 1840. In the 1940s, there were several independent announcements of a more efficient version. However, when those authors learned that their discovery was more than a century old, attempts to expand the announcement to an article seemed to have been set aside. A few articles were published that gave further explorations of the method. The expositions in Alton S. Householder, *The Theory of Matrices in Numerical Analysis*, Dover, New York, 1975. (ISBN 0-486-61781-5) and J. S. Frame, "Matrix functions and applications", *IEEE Spectrum* **1** (1964) (five articles in numbers 3–7) have good lists of published rediscoveries.

We will find that Frame's proof that the algorithm is correct reverses history and includes easy proofs of the important results cited in the historical development, but it can appear unmotivated in its drive for efficiency, so we begin with a historical treatment.

2. Powers and traces We begin by collecting some facts that are easily verified in the case when the characteristic polynomial has no repeated roots. The general case is also true, but the proof involves technical difficulties that are avoided in the modern proof.

The key to the method is that the sum of the elements on the diagonal of a square matrix M , called the **trace** of M is also equal to the sum of the eigenvalues of M . By itself, this isn't too useful. However, we also know that the eigenvalues of M^k are the k -th powers of the eigenvalues of M . Thus, computing the traces of powers of M gives the sums of powers of the roots of the characteristic polynomial.

A small example will be useful. Let

$$M = \begin{bmatrix} 3 & 1 & 5 \\ 3 & 3 & 1 \\ 4 & 6 & 4 \end{bmatrix}.$$

Then,

$$M^2 = \begin{bmatrix} 32 & 36 & 36 \\ 22 & 18 & 22 \\ 46 & 46 & 42 \end{bmatrix} \quad M^3 = \begin{bmatrix} 348 & 356 & 340 \\ 208 & 208 & 216 \\ 444 & 436 & 444 \end{bmatrix}.$$

If the eigenvalues of M are denoted $\lambda_1, \lambda_2, \lambda_3$, then equating traces and power sums gives

$$\begin{aligned} \lambda_1 + \lambda_2 + \lambda_3 &= 10 \\ \lambda_1^2 + \lambda_2^2 + \lambda_3^2 &= 92 \\ \lambda_1^3 + \lambda_2^3 + \lambda_3^3 &= 1000 \end{aligned}$$

Since M is recognized as 10 times a **Markov matrix**, we have $\lambda_1 = 10$, and without trying too hard, one finds $\lambda_3 = -\lambda_2$ and $2\lambda_2^2 = -8$. Thus, $\lambda_2 = 2i$. This gives the characteristic polynomial to be $\lambda^3 - 10\lambda^2 + 4\lambda - 40 = (\lambda - 10)(\lambda^2 + 4)$ — **if the calculations are correct**. Since we have the characteristic polynomial, we can verify the other eigenvalues of the Markov matrix $(1/10)M$ are of absolute value less than 1, although, in this case they are purely imaginary.

3. The Newton identities The coefficients of a polynomial are **elementary symmetric polynomials** of the roots found by expanding the product $\prod(\lambda - \lambda_i)$. There are relations between these elementary symmetric polynomials and the power sums used in our example that were put in a systematic form by Newton. The expressions take slightly different forms depending on how one chooses the signs of the coefficients. We choose to emphasize the coefficients of the polynomial rather than the symmetric functions. To be consistent with Faddeev, we write the characteristic polynomial in the form

$$\lambda^n - p_1\lambda^{n-1} - p_2\lambda^{n-2} - \cdots - p_n,$$

with negative signs in front of all terms except the leading term, which has coefficient +1. In particular, the determinant turns out to be $(-1)^{n-1}p_n$. We also introduce s_k to stand for the sum of the k -th powers of the roots. The first few formulas are

$$\begin{aligned} p_1 &= s_1 \\ 2p_2 &= s_2 - p_1s_1 \\ 3p_3 &= s_3 - p_1s_2 - p_2s_1 \end{aligned}$$

and, for each k ,

$$kp_k = s_k - \sum \{ p_i s_j : i > 0, j > 0, i + j = k \}$$

If these identities are used in order, then everything needed on the right side of an identity will be either a given s_n or a p_k obtained at a previous step. For now, we will accept these identities on good authority, but they will follow from Frame's proof.

4. The Twentieth Century Improvement Instead of computing the powers of M whose traces give the s_k and applying Newton's identities to these numbers, we construct the p_k and a sequence of matrices $F_k(M)$ such that

$$F_1(M) = M \quad (B)$$

$$p_k = \frac{1}{k} \operatorname{tr} (F_k(M)) \quad (T_k)$$

$$F_{k+1}(M) = M(F_k(M) - p_k I) \quad (M_k)$$

Here, (B) is the basis of the inductive construction. It tells us what $F_1(M)$ is. Then for $k = 1, 2, 3, \dots$ we apply first (T_k) to use $F_k(M)$ to identify p_k and then (M_k) to determine $F_{k+1}(M)$. The only part of the algorithm that changes with k is the division by k in (T_k). This inductive definition is used to assure that

$$F_k(M) = M^k - p_1 M^{k-1} - p_2 M^{k-2} - \dots - p_{k-1} M.$$

The Newton identities then show that (T_k) gives the correct value of p_k .

One nice feature of this algorithm is that it is self-checking. The Cayley-Hamilton theorem (mentioned in our text only in exercises — see the index) says that a matrix satisfies its characteristic polynomial. In the context of this method, this says that $F_{n+1}(M)$ must turn out to be identically zero. **If you fail to obtain a zero matrix at this stage, then you know that you have made a mistake in your calculation.** However, the operations in the algorithm consist only of modifying the diagonal of a matrix and multiplying by a fixed matrix. Frame gave a direct proof that $F_{n+1}(M)$ is zero, which will appear later in these notes, so that the Cayley-Hamilton Theorem appears as a consequence. Thus, although designed as a computational method, the construction has the theoretical benefit of including proofs of both the Newton identities and the Cayley-Hamilton theorem.

5. Revisiting the first example Here is how the previous example looks when the newer method is used.

$$F_1(M) = \begin{bmatrix} \mathbf{3} & 1 & 5 \\ 3 & \mathbf{3} & 1 \\ 4 & 6 & \mathbf{4} \end{bmatrix} \quad (B)$$

$$p_1 = (3 + 3 + 4)/1 = 10 \quad (T_1)$$

$$F_2(M) = \begin{bmatrix} \mathbf{3} & 1 & 5 \\ 3 & \mathbf{3} & 1 \\ 4 & 6 & \mathbf{4} \end{bmatrix} \begin{bmatrix} \mathbf{-7} & 1 & 5 \\ 3 & \mathbf{-7} & 1 \\ 4 & 6 & \mathbf{-6} \end{bmatrix} = \begin{bmatrix} \mathbf{2} & 26 & -14 \\ -8 & \mathbf{-12} & 12 \\ 6 & -14 & \mathbf{2} \end{bmatrix} \quad (M_1)$$

$$p_2 = (2 - 12 + 2)/2 = -4 \quad (T_2)$$

$$F_3(M) = \begin{bmatrix} \mathbf{3} & 1 & 5 \\ 3 & \mathbf{3} & 1 \\ 4 & 6 & \mathbf{4} \end{bmatrix} \begin{bmatrix} \mathbf{6} & 26 & -14 \\ -8 & \mathbf{-8} & 12 \\ 6 & -14 & \mathbf{6} \end{bmatrix} = \begin{bmatrix} \mathbf{40} & 0 & 0 \\ 0 & \mathbf{40} & 0 \\ 0 & 0 & \mathbf{40} \end{bmatrix} \quad (M_2)$$

$$p_3 = (40 + 40 + 40)/3 = 40 \quad (T_3)$$

and $F_4(M) = 0$, as expected.

A convenient way to organize the work for hand computation is to arrange the n matrices $F_k(M)$ as the **second** row of a table, using the first row for the index k . The third row should contain the p_k , including the computation shown in the lines (T_k). A fourth row should contain the matrices $F_k(M) - p_k I$ in order to have them available to use as the second factor in the computation (M_{k+1}). While our description has emphasized

the $F_k(M)$, the matrices in this fourth row turn out to have the most important applications. In particular, if M is invertible, $F_n(M)$ will be $\det(M)I$, so the previous matrix in the fourth row must be the **adjugate** of M . On the other hand, if M is singular, $F_n(M) = \mathbf{0}$, so the columns of the previous matrix in the fourth row belong to the **nullspace** of M . If it is not itself a zero matrix, its nonzero columns give eigenvectors of M for $\lambda = 0$. If it is the zero matrix, then the preceding matrix must have columns belonging to the nullspace of M . Continuing in this fashion, the last nonzero matrix has columns belonging to the nullspace of M .

6. Exercises In order to check that you have organized the steps of this algorithm correctly, begin with some examples where you know the characteristic polynomial

1. A 3 by 3 identity matrix.
2. The 4 by 4 matrix

$$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Then try a more general 4 by 4 matrix.

3. The matrix

$$\begin{bmatrix} 3 & 1 & 5 & -2 \\ 3 & 3 & 0 & 1 \\ 4 & 6 & -4 & 3 \\ 2 & -1 & -2 & 0 \end{bmatrix}$$

7. A direct proof The appearance of the adjugate in the fourth row of our computational scheme becomes a **central feature** of the method in Frame's approach. This approach identifies the bottom row of our table with coefficient matrices when the adjugate matrix of $\lambda I - M$ is written as a polynomial in λ with matrix coefficients.

The entries of the adjugate of $\lambda I - M$ are $(n - 1)$ by $(n - 1)$ subdeterminants of $\lambda I - M$, so each entry is a polynomial of degree at most $n - 1$. Alternatively, it is a polynomial of degree $n - 1$ with coefficients that are constant n by n matrices. Consistent with our previous convention, we write the adjugate of $\lambda I - M$ as

$$-\sum_{j=0}^{n-1} \lambda^{n-1-j} B_j$$

with $B_0 = -I$. Indices of summation may be suppressed by adopting the convention that j ranges over **all** integers, with the convention that $B_j = \mathbf{0}$ unless $0 \leq j \leq n - 1$. In particular, $B_n = \mathbf{0}$ **by definition**.

Since this matrix is the adjugate of $\lambda I - M$,

$$(\lambda I - M) \left(\sum \lambda^{n-1-j} B_j \right) = f(\lambda) I$$

where $f(\lambda)$ is the characteristic polynomial of M . Equating the coefficients of λ^{n-k} on each side of this equation gives

$$B_k - MB_{k-1} = -p_k I. \quad (*)$$

This shows that B_k is obtained from B_{k-1} by multiplying by M and subtracting a scalar matrix — exactly the process that we just described. To prove the algorithm correct **and** to prove the Cayley-Hamilton theorem (since we have agreed that $B_n = \mathbf{0}$), it suffices to verify that the trace of MB_{k-1} is kp_k for all k . This uses a more general result.

Lemma. If $M(x)$ is any matrix function, then

$$\frac{d}{dx} \det(M(x)) = \text{tr} \left(\text{adj}(M) \frac{d}{dx} M(x) \right)$$

Proof. Consider $\det(M(x))$ as a function of its entries $m_{ij}(x)$ and apply the chain rule to get

$$\frac{d}{dx} \det(M(x)) = \sum_{i,j} \frac{\partial}{\partial m_{ij}} \det(M(x)) \frac{d}{dx} m_{ij}(x).$$

The sum ranges over entries in the matrix. However, the partial derivatives of the determinant with respect to an entry m_{ij} is just the **cofactor** of that entry. This is written in the (j, i) position in the adjugate, so the sum over i gives the (j, j) entry of the product in the statement of the lemma. Summing over j gives the trace of this matrix, as claimed.

To apply to the characteristic polynomial, note that the derivative of $\lambda I - M$ with respect to λ is the identity matrix. Thus, the lemma shows that the coefficient of $-x^{n-k-1}$ (the minus sign appears because of our convention about naming coefficients) in the derivative of the characteristic polynomial is

$$(k - n)p_k = \text{tr } B_k$$

and (*) gives

$$\text{tr } B_k = \text{tr}(M B_{k-1}) - n p_k.$$

Combining these gives the expression for p_k in terms of B_{k-1} that was used in the algorithm.

8. Nullspaces If zero is an eigenvalue of M , the constant term of the characteristic polynomial, $-p_n$, is zero, so $M B_{n-1} = \mathbf{0}$. Thus, the column space of B_{n-1} is contained in the nullspace of M .

If $\lambda = 0$ is a **simple root** of the characteristic polynomial (i.e., $\lambda = 0$ has **algebraic multiplicity** 1), then $p_{n-1} \neq 0$. In this case $B_{n-1} = M B_{n-2} - p_{n-1} I$, so for every v in the nullspace of M , $B_{n-1} v = -p_{n-1} v$, so v lies in the column space of B_{n-1} . Thus, the nullspace of M is contained in the column space of B_{n-1} .

Combining these results, we find that these two spaces are the same. Moreover, we have $B_{n-1} \neq \mathbf{0}$, which implies that M has rank $n - 1$, so that the nullspace of M is one dimensional (i.e., the **geometric multiplicity** of $\lambda = 0$ is 1). In particular, any nonzero column of B_{n-1} generates the nullspace of M .

It is also useful to note that all B_k are polynomials in M , so that $M B_k = B_k M$ for all k .

If the algebraic multiplicity of $\lambda = 0$ is $m > 1$, then the last m coefficients of the characteristic polynomial are zero. Thus, for $k > n - m$, $p_k = 0$ and $B_k = M B_{k-1}$. It follows that $M^m B_{n-m} = B_n = \mathbf{0}$. Thus, the column space of B_{n-m} is contained in the nullspace of M^m . The nullspace of M itself is not easily characterized if the algebraic multiplicity is bigger than 1: all of its properties are obtained from the action of M on this larger space.

Meanwhile B_{n-m} is a polynomial in K with a nonzero constant term. The Euclidean algorithm for polynomials leads to another polynomial in M , $K(M)$, whose product with B_{n-m} (on either side) is the identity plus a multiple of M^m . Now, if v is in the nullspace of M^m , $B_{n-m} K(M) v = v$. The components of $K(M) v$ give the coefficients for expressing v in terms of the columns of B_{n-m} . Thus, the nullspace of M^m is contained in the column space of B_{n-m} .

Thus, the nullspace of M^m and the column space of B_{n-m} are the same. The dimension of this space is m .

9. Eigenvectors Another use of the sequence of coefficients B_k is the **synthetic division** algorithm for finding the coefficients of a polynomial $P(\lambda)$ given in terms of powers of λ as a polynomial given in terms of powers of $\lambda - c$. The constant term of such an expression is found by dividing the polynomial by writing

$$P(\lambda) = (\lambda - c) \cdot Q(\lambda) + R$$

where R is a constant. The familiar **long division** algorithm could be used to obtain $Q(\lambda)$ as a polynomial written in powers of λ , continuing as long as terms of positive degree remain. For polynomials, this process can be described as

$$a_k \lambda^k = (\lambda - c) \cdot a_k \lambda^{k-1} + ca_k \lambda^{k-1}$$

This allows a neater way to organize the computation. The first term on the right contributes to the quotient, while the second term is **added to** the original term in λ^{k-1} to get the $a_{k-1} \lambda^{k-1}$ needed to continue the process. The resulting sequence a_k are the coefficients of $Q(\lambda)$ **in powers of** λ , with the last term being the remainder, which is the constant term when $P(\lambda)$ is written in terms of powers of $(\lambda - c)$. This process is written in only three lines. The top line contains the original coefficients, from highest degree to lowest. The second line contains ca_k in the column of coefficients of λ^{k-1} , and the third line contains the sum of the first two lines. At the end of the process, the third line contains the coefficients of $Q(\lambda)$ and the remainder that will be the constant term of the new expression.

To obtain more terms in the expression, $Q(\lambda)$ should be divided by $(\lambda - c)$ in the same way.

This can be applied to matrix polynomials in exactly the same way. To find eigenvectors, it is only necessary to find **the last nonzero term** when expressed in powers of $(\lambda - c)$. If the eigenvalue is simple, this will be the constant term found by one application of synthetic division. The rule is to add c times the **new** B_{k-1} to B_k , for $k = 1, 2, 3, \dots$ to obtain the new B_k .

To return to our example, the original sequence B_0, B_1, B_2 is

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} -7 & 1 & 5 \\ 3 & -7 & 1 \\ 4 & 6 & -6 \end{bmatrix} \begin{bmatrix} 6 & 26 & -14 \\ -8 & -8 & 12 \\ 6 & -14 & 6 \end{bmatrix}.$$

When expressed in terms of powers of $(\lambda - 10)$, this becomes

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 3 & 1 & 5 \\ 3 & 3 & 1 \\ 4 & 6 & 4 \end{bmatrix} \begin{bmatrix} 36 & 36 & 36 \\ 22 & 22 & 22 \\ 46 & 46 & 46 \end{bmatrix}.$$

The last of these has columns that are eigenvectors, as predicted.

The same calculation can be done to translate by $2i$. The details a little messier, but the columns of the last matrix turn out to be multiples of $[1 + i, -i, -1]^T$.

This calculation can be done one column at a time. Most of the time, all columns are nonzero, so the first column is all that is needed to find an eigenvector. In the case of algebraic multiplicity 1, which also happens most of the time, the eigenspace is one dimensional, so a single eigenvector is all that is required.

An upper triangular matrix gives a useful way to illustrate the theoretical aspects of this calculation. All matrices appearing are triangular — including the matrix of eigenvectors. In particular, the eigenvector for the **last** eigenvalue (i.e., the one in the **last** column) has a nonzero entry in its last position. An upper triangular matrix whose columns are multiples of this vector will have all but its last column zero. In this case, one should **immediately** restrict attention to the last column in the search for an eigenvector.

End of Supplement