Numerical Analysis Hilary Term 2026
Lectures 7–8: Computing eigenvalues: The Symmetric QR Algorithm

Direct vs. Iterative Methods: methods such as LU or QR factorisations and solving $Ax = b$ using them are *direct*: they compute a certain number of operations and then finish with "the answer". Another class of methods are **iterative**:
- construct a sequence;
- truncate that sequence "after convergence";
- typically concerned with fast convergence rate (rather than operation count).

Note that unlike LU, QR or linear systems $Ax = b$, algorithms for eigenvalues are necessarily iterative: By Galois theory, no finite algorithm can compute eigenvalues of $n \times n (\geq 5)$ matrices exactly in a finite number of operations. We still have an incredibly reliable algorithm to compute them, essentially to full accuracy (for symmetric matrices; for nonsymmetric matrices, in a "backward stable" manner; this is outside the scope).

Notation: for $x \in \mathbb{R}^n$, $\|x\| = \sqrt{x^{\mathrm{T}}x}$ is the (Euclidean) length of x.

Notation: in iterative methods, x_k usually means the vector x at the kth iteration (rather than kth entry of vector x). Some sources use x^k or $x^{(k)}$ instead.

Power Iteration: a simple method for calculating a single (largest) eigenvalue of a square matrix A (and its associated eigenvector). For arbitrary $y \in \mathbb{R}^n$, set $x_0 = y/\|y\|$ to calculate an initial vector, and then for $k = 0, 1, \ldots$

Compute $y_k = Ax_k$

and set $x_{k+1} = y_k/\|y_k\|$.

This is the **Power Method** or **Power Iteration**, and computes unit vectors in the direction of $x_0, Ax_0, A^2x_0, A^3x_0, \ldots, A^kx_0$.

Suppose that A is diagonalizable so that there is a basis of eigenvectors of A:

$$\{v_1, v_2, \ldots, v_n\}$$

with $Av_i = \lambda_i v_i$ and $\|v_i\| = 1$, $i = 1, 2, \ldots, n$, and assume that

$$|\lambda_1| > |\lambda_2| \geq \cdots \geq |\lambda_n|.$$

Then we can write

$$x_0 = \sum_{i=1}^{n} \alpha_i v_i$$

for some $\alpha_i \in \mathbb{R}$, $i = 1, 2, \ldots, n$, so

$$A^k x_0 = A^k \sum_{i=1}^{n} \alpha_i v_i = \sum_{i=1}^{n} \alpha_i A^k v_i.$$

However, since $Av_i = \lambda_i v_i \implies A^2 v_i = A(Av_i) = \lambda_i Av_i = \lambda_i^2 v_i$, inductively $A^k v_i = \lambda_i^k v_i$. So

$$A^k x_0 = \sum_{i=1}^{n} \alpha_i \lambda_i^k v_i = \lambda_1^k \left[\alpha_1 v_1 + \sum_{i=2}^{n} \alpha_i \left(\frac{\lambda_i}{\lambda_1} \right)^k v_i \right].$$

Since $(\lambda_i/\lambda_1)^k \to 0$ as $k \to \infty$, $A^k x_0$ tends to look like $\lambda_1^k \alpha_1 v_1$ as k gets large. The result is that by normalizing to be a unit vector

$$\frac{A^k x_0}{\|A^k x_0\|} \to \pm v_1 \quad \text{and} \quad \frac{\|A^k x_0\|}{\|A^{k-1} x_0\|} \approx \left| \frac{\lambda_1^k \alpha_1}{\lambda_1^{k-1} \alpha_1} \right| = |\lambda_1|$$

as $k \to \infty$, and the sign of λ_1 is identified by looking at, e.g., $(A^k x_0)_1/(A^{k-1} x_0)_1$.

Essentially the same argument works when we normalize at each step: the Power Iteration may be seen to compute $y_k = \beta_k A^k x_0$ for some β_k. Then, from the above,

$$x_{k+1} = \frac{y_k}{\|y_k\|} = \frac{\beta_k}{|\beta_k|} \cdot \frac{A^k x_0}{\|A^k x_0\|} \to \pm v_1.$$

Similarly, $y_{k-1} = \beta_{k-1} A^{k-1} x_0$ for some β_{k-1}. Thus

$$x_k = \frac{\beta_{k-1}}{|\beta_{k-1}|} \cdot \frac{A^{k-1} x_0}{\|A^{k-1} x_0\|} \quad \text{and hence} \quad y_k = A x_k = \frac{\beta_{k-1}}{|\beta_{k-1}|} \cdot \frac{A^k x_0}{\|A^{k-1} x_0\|}.$$

Therefore, as above,

$$\|y_k\| = \frac{\|A^k x_0\|}{\|A^{k-1} x_0\|} \approx |\lambda_1|,$$

and the sign of λ_1 may be identified by looking at, e.g., $(x_{k+1})_1/(x_k)_1$.

Hence the largest eigenvalue (and its eigenvector) can be found.

Note: it is unlikely but possible for a chosen vector x_0 that $\alpha_1 = 0$, but rounding errors in the computation generally introduce a small component in v_1, so that in practice this is not a concern!

This simplified method for eigenvalue computation is the basis for effective methods, but the current state of the art is the **QR Algorithm** which was invented by John Francis in London in 1959/60. As we shall see, the mechanics of QR algorithm is very much related to the power method.

The QR algorithm: We now describe the QR algorithm, a magical algorithm that can solve eigenvalue problems $Ax = \lambda x$ and finds all eigenvalues (and eigenvectors).

For simplicity we consider the algorithm only in the case when A is symmetric, but it is applicable also to nonsymmetric matrices with minor modifications.

Recall: a symmetric matrix A is similar to B if there is a nonsingular matrix P for which $A = P^{-1}BP$. Similar matrices have the same eigenvalues, since if $A = P^{-1}BP$,

$$0 = \det(A - \lambda I) = \det(P^{-1}(B - \lambda I)P) = \det(P^{-1}) \det(P) \det(B - \lambda I),$$

so $\det(A - \lambda I) = 0$ if, and only if, $\det(B - \lambda I) = 0$.

The basic **QR algorithm** is:
    Set $A_1 = A$.
    for $k = 1, 2, \ldots$
       form the QR factorization $A_k = Q_k R_k$
       and set $A_{k+1} = R_k Q_k$

```
end
```

Proposition. The matrices $A_1, A_2, \ldots, A_k, \ldots$ are all symmetric and similar, and thus have the same eigenvalues.

Proof. Since

$$A_{k+1} = R_k Q_k = (Q_k^T Q_k) R_k Q_k = Q_k^T (Q_k R_k) Q_k = Q_k^T A_k Q_k = Q_k^{-1} A_k Q_k,$$

A_{k+1} is symmetric if A_k is, and is similar to A_k. □

At least when A has eigenvalues of distinct modulus $|\lambda_1| > |\lambda_2| > \cdots > |\lambda_n|$, this basic QR algorithm can be shown to work (A_k converges to a diagonal matrix as $k \to \infty$, the diagonal entries of which are the eigenvalues). To see this, we make the following observations.

Lemma. The iterates A_k of the QR algorithm satisfy (defining $Q^{(k)} := Q_1 \cdots Q_k$ and $R^{(k)} := R_k \cdots R_1$)

$$A_{k+1} = (Q^{(k)})^T A Q^{(k)}, \tag{1}$$

and

$$A^k = (Q_1 \cdots Q_k)(R_k \cdots R_1) = Q^{(k)} R^{(k)} \tag{2}$$

is the QR factorization of A^k.

Proof. (1) follows from a repeated application of the above proposition.

We use induction for (2): $k = 1$ trivial. Suppose $A^{k-1} = Q^{(k-1)} R^{(k-1)}$. Then $A_k = R_{k-1} Q_{k-1} = (Q^{(k-1)})^T A Q^{(k-1)}$, and

$$(Q^{(k-1)})^T A Q^{(k-1)} = Q_k R_k.$$

Then $A Q^{(k-1)} = Q^{(k-1)} Q_k R_k$, and so

$$A^k = A Q^{(k-1)} R^{(k-1)} = Q^{(k-1)} Q_k R_k R^{(k-1)} = Q^{(k)} R^{(k)},$$

giving (2). □

Let us now connect the above lemma with the power method.

Lemma. With $Q^{(k)}$ as in (2), let q_1 be its first column, and let $e_1 = [1, 0, \ldots, 0]^T$. Then q_1 is equal to either $\frac{A^k e_1}{\|A^k e_1\|_2}$ or $-\frac{A^k e_1}{\|A^k e_1\|_2}$.

Proof. Right-multiplying e_1 to (2) yields $A^k e_1 = Q^{(k)} R^{(k)} e_1$. Since $R^{(k)}$ is upper triangular $R^{(k)} e_1 = [R_{1,1}^{(k)}, 0, \ldots, 0]^T$, and so $Q^{(k)} R^{(k)} e_1$ is parallel to q_1, which has unit norm. □

The results show in particular that the first column q_1 of $Q^{(k)}$ is the result of power method applied k times to the initial vector $e_1 = [1, 0, \ldots, 0]^T$. It then follows that q_1 converges to the dominant eigenvector. The second vector then starts converging to the 2nd dominant eigenvector, and so on. Once the columns of $Q^{(k)}$ converge to eigenvectors (note that they are orthogonal by design), (1) shows that A_k converge to a diagonal matrix of eigenvalues.

However, a really practical, fast algorithm is based on some refinements.

Reduction to tridiagonal form: the idea is to apply explicit similarity transformations $QAQ^{-1} = QAQ^T$, with Q orthogonal, so that QAQ^T is tridiagonal.

Note: direct reduction to triangular form would reveal the eigenvalues, but is not possible. If

$$H(w)A = \begin{bmatrix} \times & \times & \cdots & \times \\ 0 & \times & \cdots & \times \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \times & \cdots & \times \end{bmatrix}$$

then $H(w)AH(w)^{\mathrm{T}}$ is generally full, i.e., all zeros created by pre-multiplication are destroyed by the post-multiplication. However, if

$$A = \begin{bmatrix} \gamma & u^{\mathrm{T}} \\ u & C \end{bmatrix}$$

(as $A = A^{\mathrm{T}}$) and

$$w = \begin{bmatrix} 0 \\ \hat{w} \end{bmatrix} \quad \text{where} \quad H(\hat{w})u = \begin{bmatrix} \alpha \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

it follows that

$$H(w)A = \begin{bmatrix} \gamma & & u^{\mathrm{T}} & \\ \alpha & \times & \vdots & \times \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \times & \vdots & \times \end{bmatrix},$$

i.e., the u^{T} part of the first row of A is unchanged. However, then

$$H(w)AH(w)^{-1} = H(w)AH(w)^{\mathrm{T}} = H(w)AH(w) = \begin{bmatrix} \gamma & \alpha & 0 & \cdots & 0 \\ \hline \alpha & & & & \\ 0 & & & B & \\ \vdots & & & & \\ 0 & & & & \end{bmatrix},$$

where $B = H(\hat{w})CH^{\mathrm{T}}(\hat{w})$, as $u^{\mathrm{T}}H(\hat{w})^{\mathrm{T}} = (\alpha, \quad 0, \quad \cdots, \quad 0)$; note that $H(w)AH(w)^{\mathrm{T}}$ is symmetric as A is.

Now we inductively apply this to the smaller matrix B, as described for the QR factorization but using post- as well as pre-multiplications. The result of $n-2$ such Householder similarity transformations is the matrix

$$H(w_{n-2}) \cdots H(w_2)H(w)AH(w)H(w_2) \cdots H(w_{n-2}),$$

which is tridiagonal.

The QR factorization of a tridiagonal matrix can now easily be achieved with $n-1$ *Givens* rotations $J(i,j)$; these are orthogonal matrices that are I except for the four elements: the $(i,i),(i,j),(j,i),(j,j)$ entries with values $c, s, -s, c$ respectively, where $c^2 + s^2 = 1$

(cosine and sine); one can choose c s.t. $\begin{bmatrix} c & s \\ -s & c \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} \sqrt{a^2 + b^2} \\ 0 \end{bmatrix}$. (The operations below can be done with Householder matrices too, but Givens rotations are more straightforward).

Now if A is tridiagonal

$$\underbrace{J(n-1,n)\cdots J(2,3)J(1,2)}_{Q^{\mathrm{T}}}A = R, \quad \text{upper triangular.}$$

Precisely, R has a diagonal and 2 super-diagonals,

$$R = \begin{bmatrix} \times & \times & \times & 0 & 0 & 0 & \cdots & 0 \\ 0 & \times & \times & \times & 0 & 0 & \cdots & 0 \\ 0 & 0 & \times & \times & \times & 0 & \cdots & 0 \\ \vdots & \vdots & & & & & & \vdots \\ 0 & 0 & 0 & 0 & \times & \times & \times & 0 \\ 0 & 0 & 0 & 0 & 0 & \times & \times & \times \\ 0 & 0 & 0 & 0 & 0 & 0 & \times & \times \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \times \end{bmatrix}$$

(exercise: check!). In the QR algorithm, the next matrix in the sequence is RQ.

Lemma. In the QR algorithm applied to a symmetric tridiagonal matrix, the symmetry and tridiagonal form are preserved, that is, A_k is symmetric tridiagonal for all k.

Proof. We have already shown that if $A_k = QR$ is symmetric, then so is $A_{k+1} = RQ$. If $A_k = QR = J(1,2)^{\mathrm{T}}J(2,3)^{\mathrm{T}}\cdots J(n-1,n)^{\mathrm{T}}R$ is tridiagonal, then $A_{k+1} = RQ = RJ(1,2)^{\mathrm{T}}J(2,3)^{\mathrm{T}}\cdots J(n-1,n)^{\mathrm{T}}$. Recall that post-multiplication of a matrix by $J(i,i+1)^{\mathrm{T}}$ replaces columns i and $i+1$ by linear combinations of the pair of columns, while leaving columns $j = 1, 2, \ldots, i-1, i+2, \ldots, n$ alone. Thus, since R is upper triangular, the only subdiagonal entry in $RJ(1,2)^{\mathrm{T}}$ is in position $(2,1)$. Similarly, the only subdiagonal entries in $RJ(1,2)^{\mathrm{T}}J(2,3)^{\mathrm{T}} = (RJ(1,2)^{\mathrm{T}})J(2,3)^{\mathrm{T}}$ are in positions $(2,1)$ and $(3,2)$. Inductively, the only subdiagonal entries in

$$RJ(1,2)^{\mathrm{T}}J(2,3)^{\mathrm{T}}\cdots J(i-2,i-1)^{\mathrm{T}}J(i-1,i)^{\mathrm{T}}$$
$$= (RJ(1,2)^{\mathrm{T}}J(2,3)^{\mathrm{T}}\cdots J(i-2,i-1)^{\mathrm{T}})J(i-1,i)^{\mathrm{T}}$$

are in positions $(j,j-1)$, $j = 2, \ldots i$. So, the lower triangular part of A_{k+1} only has nonzeros on its first subdiagonal. However, then since A_{k+1} is symmetric, it must be tridiagonal. $\qquad\square$

Using shifts. One further and final step in making an efficient algorithm is the use of **shifts**:

    for $k = 1, 2, \ldots$
        form the QR factorization of $A_k - \mu_k I = Q_k R_k$
        and set $A_{k+1} = R_k Q_k + \mu_k I$

```
end
```

For any chosen sequence of values of $\mu_k \in \mathbb{R}$, $\{A_k\}_{k=1}^{\infty}$ are symmetric and tridiagonal if A_1 has these properties, and similar to A_1.

The simplest shift to use is $a_{n,n}$, which leads rapidly in almost all cases to

$$A_k = \left[\begin{array}{c|c} T_k & 0 \\ \hline 0^{\mathrm{T}} & \lambda \end{array} \right],$$

where T_k is $n-1$ by $n-1$ and tridiagonal, and λ is an eigenvalue of A_1. Inductively, once this form has been found, the QR algorithm with shift $a_{n-1,n-1}$ can be concentrated only on the $n-1$ by $n-1$ leading submatrix T_k. This process is called **deflation**.

Why does introducing shifts help? To understand this we establish a connection between QR and the power method applied to the *inverse* (known as the *inverse power method*).

Lemma. With $Q^{(k)}$ as in (2), denote by q_n its last column, and let $e_n = [0, 0, \ldots, 1]^T$. Then q_n is equal to either $\frac{A^{-k}e_n}{\|A^{-k}e_1\|_2}$ or $-\frac{A^{-k}e_n}{\|A^{-k}e_1\|_2}$.

Proof. Recall (2), and take the inverse:

$$A^{-k} = (R^{(k)})^{-1}(Q^{(k)})^T,$$

and take the transpose:

$$(A^{-k})^T (= A^{-k}) = Q^{(k)}(R^{(k)})^{-T}.$$

Now multiplying e_n gives

$$A^{-k}e_n = Q^{(k)}(R^{(k)})^{-T}e_n.$$

Since $(R^{(k)})^{-T}$ is lower triangular, it follows that $Q^{(k)}(R^{(k)})^{-T}e_n$ is parallel to q_n. $\qquad\square$

This shows that the **final** column of $Q^{(k)}$ is the result of power method applied to e_n now with the **inverse** A^{-1}. Thus the last column of $Q^{(k)}$ is converging to the eigenvector for the smallest eigenvalue λ_n, with convergence factor $|\frac{\lambda_n}{\lambda_{n-1}}|$; $Q^{(k)}$ is converging not only from the first, but (more significantly) from the last column(s).

Now we see how the introduction of shift has a drastic effect on the convergence: it changes the factor to $|\frac{\lambda_{\sigma(n)}-\mu}{\lambda_{\sigma(n-1)}-\mu}|$, where σ is a permutation such that $|\lambda_{\sigma(1)} - \mu| \geq |\lambda_{\sigma(2)} - \mu| \geq \cdots \geq |\lambda_{\sigma(n)} - \mu|$. If μ is close to an eigenvalue, this implies (potentially extremely) fast convergence; in fact by choosing the shift $\mu_k = a_{n,n}$, it can be shown that (proof omitted and non-examinable) $a_{m,m-1}$ converges *cubically*: $|a_{m,m-1,k+1}| = O(|a_{m,m-1,k}|^3)$.

The overall algorithm for calculating the eigenvalues of an n by n symmetric matrix:

```
reduce A to tridiagonal form by orthogonal
(Householder) similarity transformations.

for m = n, n − 1, . . . 2
  while a_{m−1,m} > tol
    [Q, R] = qr(A − a_{m,m}I)
    A = RQ + a_{m,m}I
  end while
  record eigenvalue λ_m = a_{m,m}
```

```
        A ← leading m − 1 by m − 1 submatrix of A
end
record eigenvalue λ₁ = a₁,₁
```

$A \leftarrow$ `leading` $m-1$ `by` $m-1$ `submatrix of` A

`end`

`record eigenvalue` $\lambda_1 = a_{1,1}$

Computing roots of polynomials via eigenvalues Let us describe a nice application of computing eigenvalues (by the QR algorithm). Let $p(x) = \sum_{i=0}^{n} c_i x^i$ be a degree-n polynomial so that $c_n \neq 0$, and suppose we want to find its roots, i.e., values of λ for which $p(\lambda) = 0$; there are n of them in $\mathbb{C}$. For example, $p(x)$ might be an approximant to data, obtained by Lagrange interpolation from the first lecture. Why roots? For example, you might be interested in the minimum of p; this can be obtained by differentiating and setting to zero $p'(x) = 0$, which is again a polynomial rootfinding problem (for p').

How do we solve $p(x) = 0$? Recall that eigenvalues of A are the roots of its characteristic polynomial. Here we take the opposite direction—construct a matrix whose characteristic polynomial is p.

Consider the following matrix, which is called the **companion matrix** (the blank elements are all 0) for the polynomial $p(x) = \sum_{i=0}^{n} c_i x^i$:

$$
C = \begin{bmatrix}
-\frac{c_{n-1}}{c_n} & -\frac{c_{n-2}}{c_n} & \cdots & -\frac{c_1}{c_n} & -\frac{c_0}{c_n} \\
1 & & & & \\
& 1 & & & \\
& & \ddots & & \\
& & & 1 & 0
\end{bmatrix}. \tag{3}
$$

Then direct calculation shows that if $p(\lambda) = 0$ then $Cx = \lambda x$ with $x = [\lambda^{n-1}, \lambda^{n-2}, \ldots, \lambda, 1]^T$. Indeed one can show that the characteristic polynomial is $\det(\lambda I - C) = p(\lambda)/c_n$ (nonexaminable), so this implication is necessary and sufficient, so the eigenvalues of C are precisely the roots of p, counting multiplicities.

Thus to compute roots of polynomials, one can compute eigenvalues of the companion matrix via the QR algorithm—this turns out to be a very powerful idea!