Mathematical Tripos Part II: Michaelmas Term 2022

Numerical Analysis – Lecture 21

5 Eigenvalues and eigenvectors

We consider in this chapter the problem of computing eigenvalues and eigenvectors of matrices. Let A be a real $n \times n$ matrix. The eigenvalue equation is $A\boldsymbol{w} = \lambda \boldsymbol{w}$, where λ is a scalar, which may be complex in general, and $\boldsymbol{w}$ is a nonzero vector. If A is diagonalizable, then the eigenvectors form a basis of $\mathbb{R}^n$. If A is symmetric, we know that the eigenvalues are all real, and that the eigenvectors form an orthonormal basis of $\mathbb{R}^n$.

We start by describing algorithms to compute a single eigenvalue/eigenvector pair for A. In this chapter we use $\|\cdot\|$ to denote the Euclidean norm on $\mathbb{C}^n$, i.e.,

$$\|\boldsymbol{x}\|^2 = \boldsymbol{x}^*\boldsymbol{x} = \sum_{i=1}^n |x_i|^2.$$

5.1 Power method

The iterative algorithms that will be studied for the calculation of eigenvalues and eigenvectors are all closely related to the *power method*, which has the following basic form for generating a single eigenvalue and eigenvector of A. We pick a nonzero vector $\boldsymbol{x}^{(0)}$ in $\mathbb{R}^n$. Then, for $k = 0, 1, 2, \ldots$, we let $\boldsymbol{x}^{(k+1)}$ be a nonzero multiple of $A\boldsymbol{x}^{(k)}$, so that $\|\boldsymbol{x}^{(k+1)}\| = 1$.

POWER ITERATION: for $k = 0, 1, 2, \ldots$

- Set $\boldsymbol{y} = A\boldsymbol{x}^{(k)}$
- $\boldsymbol{x}^{(k+1)} = \boldsymbol{y}/\|\boldsymbol{y}\|$

The next theorem shows that the sequence $\boldsymbol{x}^{(k)}$ converges to an eigenvector of A associated with the largest eigenvalue in modulus, provided all the other eigenvalues of A have strictly smaller magnitude. Observe that the eigenvectors of A are only specified up to a scalar multiple, and for this reason, the theorem below studies the distance between $\boldsymbol{x}^{(k)}$ and the *linear span of* $\boldsymbol{w}_1$, where $\boldsymbol{w}_1$ is an eigenvector of A (instead of just the distance $\|\boldsymbol{x}^{(k)} - \boldsymbol{w}_1\|$). If $\boldsymbol{x}$ and $\boldsymbol{w}$ are two vectors in $\mathbb{C}^n$ with unit Euclidean length, then it is easy to check that

$$\operatorname{dist}(\boldsymbol{x}, \operatorname{span}(\boldsymbol{w}))^2 = \min_{\alpha \in \mathbb{C}} \|\boldsymbol{x} - \alpha \boldsymbol{w}\|^2 = 1 - |\boldsymbol{x}^*\boldsymbol{w}|^2, \tag{5.1}$$

which is attained at $\alpha = \boldsymbol{w}^*\boldsymbol{x}$.

Theorem 5.1 *Assume $A \in \mathbb{C}^{n\times n}$ is diagonalizable and that its eigenvalues can be ordered in such a way that $|\lambda_1| > |\lambda_2| \geq \cdots \geq |\lambda_n|$. Let $\boldsymbol{w}_i \in \mathbb{C}^n$ be the corresponding eigenvectors of A with $\|\boldsymbol{w}_i\| = 1$. Assume $\boldsymbol{x}^{(0)} = \sum_{i=1}^n c_i \boldsymbol{w}_i$ with $c_1 \neq 0$. Then $\operatorname{dist}(\boldsymbol{x}^{(k)}, \mathit{span}(\boldsymbol{w}_1)) = \mathcal{O}(\rho^k)$ as $k \to \infty$, where $\rho = |\lambda_2/\lambda_1| < 1$.*

Proof. Observe that $\boldsymbol{x}^{(k)}$ is a multiple of $A^k\boldsymbol{x}^{(0)}$, which according to the decomposition of $\boldsymbol{x}^{(0)}$ can be written as

$$A^k\boldsymbol{x}^{(0)} = \sum_{i=1}^n c_i\lambda_i^k \boldsymbol{w}_i = c_1\lambda_1^k\Big(\boldsymbol{w}_1 + \sum_{i=2}^n \tfrac{c_i}{c_1}\big(\tfrac{\lambda_i}{\lambda_1}\big)^k \boldsymbol{w}_i\Big) = c_1\lambda_1^k(\boldsymbol{w}_1 + \boldsymbol{v}^{(k)}) \tag{5.2}$$

where $\boldsymbol{v}^{(k)} = \sum_{i=2}^n (c_i/c_1)(\lambda_i/\lambda_1)^k \boldsymbol{w}_i$. By our assumptions we know that $\|\boldsymbol{v}^{(k)}\| \to 0$ as $k \to \infty$, more precisely $\|\boldsymbol{v}^{(k)}\| = O(\rho^k)$. Since $\|\boldsymbol{x}^{(k)}\| = 1$, and $\boldsymbol{x}^{(k)}$ is proportional to $A^k\boldsymbol{x}^{(0)}$, we can write using (5.2)

$$\boldsymbol{x}^{(k)} = s_k \frac{\boldsymbol{w}_1 + \boldsymbol{v}^{(k)}}{\|\boldsymbol{w}_1 + \boldsymbol{v}^{(k)}\|} \tag{5.3}$$

where $s_k = c_1\lambda_1^k/|c_1\lambda_1^k|$ which satisfies $|s_k| = 1$. Thus we get

$$|\boldsymbol{w}_1^*\boldsymbol{x}^{(k)}|^2 = \frac{|1+\boldsymbol{w}_1^*\boldsymbol{v}^{(k)}|^2}{\|\boldsymbol{w}_1+\boldsymbol{v}^{(k)}\|^2} = 1 + \frac{|\boldsymbol{w}_1^*\boldsymbol{v}^{(k)}|^2 - \|\boldsymbol{v}^{(k)}\|^2}{\|\boldsymbol{w}_1+\boldsymbol{v}^{(k)}\|^2} = 1 + O(\rho^{2k}).$$

From (5.1), we get $\mathrm{dist}(\boldsymbol{x}^{(k)}, \mathrm{span}(\boldsymbol{w}_1))^2 = 1 - |\boldsymbol{w}_1^*\boldsymbol{x}^{(k)}|^2 = O(\rho^{2k})$ as desired. $\square$

Rayleigh quotient The convergence theorem above shows convergence of $\boldsymbol{x}^{(k)}$ to an eigenvector of A. What about the eigenvalue? The following definition will be important for the rest of this chapter.

Definition 5.2 *The* Rayleigh quotient *of A at a nonzero vector $\boldsymbol{x} \in \mathbb{C}^n$ is defined by*

$$r(\boldsymbol{x}) = \frac{\boldsymbol{x}^*A\boldsymbol{x}}{\boldsymbol{x}^*\boldsymbol{x}}. \tag{5.4}$$

If $A\boldsymbol{x} = \lambda\boldsymbol{x}$ then clearly $r(\boldsymbol{x}) = \lambda$. In general, $r(\boldsymbol{x}) = \arg\min_{\mu\in\mathbb{C}} \|A\boldsymbol{x} - \mu\boldsymbol{x}\|_2^2$, since $\|A\boldsymbol{x} - \mu\boldsymbol{x}\|_2^2 = |\mu|^2\boldsymbol{x}^*\boldsymbol{x} - 2\mathrm{Re}\,[\bar{\mu}\boldsymbol{x}^*A\boldsymbol{x}] + \|A\boldsymbol{x}\|^2$, which is minimized precisely at $\mu = r(\boldsymbol{x})$. One can show, using the proof of the theorem above, that the sequence of Rayleigh quotients $r(\boldsymbol{x}^{(k)})$ converges to λ_1 at the rate $\mathcal{O}(\rho^k)$. Indeed, from (5.3), we have $\boldsymbol{x}^{(k)}/s_k = \frac{\boldsymbol{w}_1+\boldsymbol{v}^{(k)}}{\|\boldsymbol{w}_1+\boldsymbol{v}^{(k)}\|}$ with $|s_k| = 1$, so that $r(\boldsymbol{x}^{(k)}) = r(\boldsymbol{x}^{(k)}/s_k) = \frac{1}{\|\boldsymbol{w}_1+\boldsymbol{v}^{(k)}\|^2}(\boldsymbol{w}_1+\boldsymbol{v}^{(k)})^*A(\boldsymbol{w}_1+\boldsymbol{v}^{(k)}) \to \lambda_1$ at the rate $O(\rho^k)$.

Deficiencies of the power method The power method may perform adequately if $c_1 \neq 0$ and $|\lambda_1| > |\lambda_2|$, but often it is unacceptably slow. Moreover, $|\lambda_1| = |\lambda_2|$ is not uncommon when A is real and nonsymmetric, because the spectral radius of A may be due to a complex conjugate pair of eigenvalues. Next, we will study the inverse iteration with shifts, which will allow us to speed up the convergence of the power method.

5.2 Inverse iteration

Inverse iteration is the power method applied to the matrix $(A - sI)^{-1}$, for some *shift* $s \in \mathbb{R}$. The eigenvalues of $(A-sI)^{-1}$ are equal to $\frac{1}{\lambda_i - s}$ where λ_i are the eigenvalues of A, and the eigenvectors are the same as those of A. Let λ be the eigenvalue of A closest to s, and let λ' be the eigenvalue second-closest to s, so that $|\lambda - s| < |\lambda' - s|$. Then, from the analysis of the power method, we know that inverse iteration will converge to an eigenvector of λ with rate ρ^k, where $\rho = \frac{|\lambda - s|}{|\lambda' - s|} < 1$.

INVERSE ITERATION: for $k = 0, 1, 2, \ldots$

- Solve $(A - sI)\boldsymbol{y} = \boldsymbol{x}^{(k)}$ (in $\boldsymbol{y}$, using e.g., LU decomposition)
- $\boldsymbol{x}^{(k+1)} = \boldsymbol{y}/\|\boldsymbol{y}\|$

The advantage of inverse iteration is the choice of the parameter s: if we have a good estimate of the eigenvalue λ, then the iterations converge very fast.

5.3 Rayleigh quotient iteration

A natural estimate for the eigenvalue λ at iteration k is the Rayleigh quotient $r(\boldsymbol{x}^{(k)})$. In Rayleigh quotient iteration, we update the shift at each iteration by the Rayleigh quotient, namely:

RAYLEIGH QUOTIENT ITERATION: for $k = 0, 1, 2, \ldots$

- $s_k = r(\boldsymbol{x}^{(k)})$
- Solve $(A - s_kI)\boldsymbol{y} = \boldsymbol{x}^{(k)}$
- $\boldsymbol{x}^{(k+1)} = \boldsymbol{y}/\|\boldsymbol{y}\|$

In practice, the convergence of Rayleigh quotient iteration is *extremely fast*.

Example 5.3 *Consider the matrix*

$$A = \begin{bmatrix} 2 & -1 & & \\ -1 & 2 & \ddots & \\ & \ddots & \ddots & -1 \\ & & -1 & 2 \end{bmatrix}$$

with $n = 5$, and the initial vector $\boldsymbol{x}^{(0)} = (1, \ldots, 1)/\sqrt{5}$. We know that the eigenvalues of A are equal to $4\sin^2(\ell\pi/(2(n+1)))$, $\ell = 1, \ldots, n$, and that the eigenvectors correspond to sinusoidal vectors with frequencies $\ell = 1, \ldots, n$. The initial vector $\boldsymbol{x}^{(0)}$ here is constant, so it makes sense to think that the Rayley quotient iteration will converge to the eigenvalue corresponding to the smallest frequency, i.e., $\ell = 1$, which in this case is $4\sin^2(\pi/12) \approx 0.267949192431$. After 3 iterations of Rayleigh quotient iteration we obtain the approximation 0.267949192649 which is correct up to 9 digits!