

24.4 The Power Method

Chapter 19 computed a Perron vector three times and each time stopped at the same place. Section 4 bracketed $\rho(\mathbf{A})$ by the ratios of $\mathbf{A}\mathbf{x}$ to $\mathbf{x}$ and observed: “the iteration $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ is the power method, and Chapter 24 studies how fast it converges.” Section 8 said the same of PageRank — “Chapter 24 studies it for general matrices, where the rate is governed by the ratio of the two largest eigenvalue moduli” — and Section 9’s ledger listed “the power method, the algorithm that computes a Perron vector at the size of the web” among the debts. Chapter 17 §01 and Chapter 20 §06 each promised, in the same words, that “Chapter 24 builds an iteration on” the Rayleigh quotient.

This section pays all five. It is about one idea used three times: **multiply by a matrix repeatedly and the largest eigenvalue takes over**. Applied to $\mathbf{A}$ itself, the idea finds the eigenvalue of largest modulus. Applied to $(\mathbf{A} - \mu\mathbf{I})^{-1}$, it finds the eigenvalue nearest μ — any eigenvalue we choose to aim at. Applied to $(\mathbf{A} - \mu\mathbf{I})^{-1}$ with μ refreshed at every step from the Rayleigh quotient, it converges cubically, and that is the iteration Chapters 17 and 20 promised.

Throughout, $F = \mathbb{R}$ or $F = \mathbb{C}$, matrices are in $M_n(F)$ read inside $M_n(\mathbb{C})$, and $\|\cdot\| = \|\cdot\|_2$. Iteration indices are superscripts in parentheses, so that subscripts stay free for coordinates.

24.4.1 Multiplying by the matrix

Take $\mathbf{A}$ diagonalizable with eigenbasis $\mathbf{v}_1, \dots, \mathbf{v}_n$ and eigenvalues $\lambda_1, \dots, \lambda_n$, and expand a starting vector:

$$\mathbf{x} = c_1\mathbf{v}_1 + \dots + c_n\mathbf{v}_n \implies \mathbf{A}^k\mathbf{x} = c_1\lambda_1^k\mathbf{v}_1 + \dots + c_n\lambda_n^k\mathbf{v}_n.$$

Divide by λ_1^k , and suppose $|\lambda_1|$ is strictly the largest. Every coefficient except the first is multiplied by $(\lambda_i/\lambda_1)^k$, a number of modulus less than 1 raised to the k -th power. So the sum collapses onto its first term. Nothing here needs to be discovered; it needs to be named, and its rate measured.

The power method is: multiply, rescale, repeat, and read off the dominant eigenvector.

Definition 24.4.1. Power Iteration

Let $\mathbf{A} \in M_n(F)$ and let $\mathbf{x}^{(0)} \in F^n$ with $\|\mathbf{x}^{(0)}\| = 1$. The **power iteration** generated by $\mathbf{A}$ and $\mathbf{x}^{(0)}$ is the sequence defined for $k \geq 0$ by

$$\mathbf{x}^{(k+1)} = \frac{\mathbf{A}\mathbf{x}^{(k)}}{\|\mathbf{A}\mathbf{x}^{(k)}\|},$$

which is defined **as long as** $\mathbf{A}\mathbf{x}^{(k)} \neq \mathbf{0}$, together with the **eigenvalue estimates**

$$\mu_k = (\mathbf{x}^{(k)})^* \mathbf{A} \mathbf{x}^{(k)}.$$

Clause by clause. The vectors are **unit** vectors: the normalization has no mathematical content, since $\mathbf{x}^{(k)}$ spans the same line as $\mathbf{A}^k\mathbf{x}^{(0)}$, and its only purpose is to stop the entries from growing or shrinking like λ_1^k until the arithmetic can no longer hold them. The iteration is **partial**: if $\mathbf{A}\mathbf{x}^{(k)} = \mathbf{0}$ the sequence stops, and it stops exactly when $\mathbf{x}^{(k)} \in \text{nul}(\mathbf{A})$ — which, as we shall see, a good starting vector avoids. The estimate μ_k is the natural number to attach to a unit

vector $\mathbf{x}$ that is nearly an eigenvector: if $\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$ exactly then $\mathbf{x}^* \mathbf{A} \mathbf{x} = \lambda$. For Hermitian $\mathbf{A}$ it is the Rayleigh quotient $R_{\mathbf{A}}(\mathbf{x}^{(k)})$ of Definition 17.1.1, whose denominator is 1 here; for a general $\mathbf{A}$ the book reserves that name for the Hermitian case, so we write the quadratic form out.

Examples.

- **The identity.** $\mathbf{A} = \mathbf{I}$: every unit vector is fixed, $\mathbf{x}^{(k)} = \mathbf{x}^{(0)}$ for all k , and $\mu_k = 1$. The degenerate case is worth a glance, because it shows that the method converges to *an* eigenvector, never to a distinguished one.
- **A diagonal matrix.** $\mathbf{A} = \text{diag}(3, 1)$ and $\mathbf{x}^{(0)} = \frac{1}{\sqrt{2}}(1, 1)$: then $\mathbf{A}^k \mathbf{x}^{(0)} \propto (3^k, 1)$, and the normalized vector runs to $\mathbf{e}_1$ with second coordinate 3^{-k} . The rate 3^{-k} is $|\lambda_2/\lambda_1|^k$.
- **A stochastic matrix.** For the Google matrix $\mathbf{G}$ of Chapter 19 §08, $\lambda_1 = 1$ and every other eigenvalue has modulus at most α (Theorem 19.8.2), so the iteration $\mathbf{x}_{k+1} = \mathbf{G}\mathbf{x}_k$ of Proposition 19.8.3 is a power iteration with rate at most α . Chapter 19 proved the rate directly from stochasticity, with no hypothesis on $\mathbf{G}$ beyond that; the theorem below proves it from the eigenvalues, for any diagonalizable matrix.

Nonexample by minimal change. Replace $\text{diag}(3, 1)$ by $\text{diag}(3, -3)$. The eigenvalues now have **equal modulus**, the collapse never happens, and from $\mathbf{x}^{(0)} = \frac{1}{\sqrt{2}}(1, 1)$ the iterates alternate between $\frac{1}{\sqrt{2}}(1, 1)$ and $\frac{1}{\sqrt{2}}(1, -1)$ forever. One clause of the hypothesis has been dropped — strictness in $|\lambda_1| > |\lambda_2|$ — and the conclusion is gone with it.

Here is the theorem, with the rate.

Theorem 24.4.1. The Power Method

Let $\mathbf{A} \in M_n(F)$ be diagonalizable, with a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of eigenvectors, $\mathbf{A}\mathbf{v}_i = \lambda_i \mathbf{v}_i$, ordered by **modulus** so that

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

Let $\mathbf{x}^{(0)} = c_1 \mathbf{v}_1 + \dots + c_n \mathbf{v}_n$ be a unit vector with $c_1 \neq 0$, and put $\rho = |\lambda_2/\lambda_1| < 1$, $\mathbf{q}_1 = \mathbf{v}_1/\|\mathbf{v}_1\|$ and

$$C = \frac{1}{\|\mathbf{v}_1\|} \sum_{i \geq 2} \left| \frac{c_i}{c_1} \right| \|\mathbf{v}_i\|.$$

Then the power iteration of Definition 24.4.1 is defined for every k , and for every k with $C\rho^k \leq \frac{1}{2}$ there is a scalar ω_k of modulus 1 with

- (a) $\|\omega_k \mathbf{x}^{(k)} - \mathbf{q}_1\| \leq 4C\rho^k$;
- (b) $|\mu_k - \lambda_1| \leq 8\|\mathbf{A}\| C\rho^k$.

In particular $\mu_k \rightarrow \lambda_1$, and the iterates approach the eigenline $\text{span}(\mathbf{v}_1)$ in the sense of (a), both at the rate ρ^k .

Idea. The whole content is the display at the head of this subsection: divide $\mathbf{A}^k \mathbf{x}^{(0)}$ by $c_1 \lambda_1^k$ and what is left is $\mathbf{v}_1$ plus a tail of size ρ^k . Two pieces of bookkeeping then stand between that and the statement. First, normalizing: a bound on $\mathbf{y} - \mathbf{v}_1$ is not yet a bound on $\mathbf{y}/\|\mathbf{y}\| - \mathbf{v}_1/\|\mathbf{v}_1\|$, and the conversion costs a factor 2 and the hypothesis $C\rho^k \leq \frac{1}{2}$, which only says that we are past the point where the tail could swamp $\mathbf{v}_1$. Second, phases: over $\mathbb{C}$ the scalar $c_1 \lambda_1^k$ has a modulus and an argument, and normalizing kills the modulus but not the argument. That argument is the ω_k . It cancels out of μ_k , because a quadratic form does

not see it.

Proof. **The iterates are defined, and what they are.** Since $c_1\lambda_1^k \neq 0$, the vector $\mathbf{A}^k \mathbf{x}^{(0)} = \sum_i c_i \lambda_i^k \mathbf{v}_i$ has a nonzero coefficient on $\mathbf{v}_1$, so $\mathbf{A}^k \mathbf{x}^{(0)} \neq \mathbf{0}$ for every k , the eigenvectors being a basis. An induction on k then gives $\mathbf{x}^{(k)} = \mathbf{A}^k \mathbf{x}^{(0)} / \|\mathbf{A}^k \mathbf{x}^{(0)}\|$: true at $k = 0$, and if it holds at k then $\mathbf{A} \mathbf{x}^{(k)}$ is a positive multiple of $\mathbf{A}^{k+1} \mathbf{x}^{(0)} \neq \mathbf{0}$, so normalizing it gives the same unit vector as normalizing $\mathbf{A}^{k+1} \mathbf{x}^{(0)}$.

The tail. Put

$$\mathbf{y}_k = \frac{\mathbf{A}^k \mathbf{x}^{(0)}}{c_1 \lambda_1^k} = \mathbf{v}_1 + \mathbf{e}_k, \quad \mathbf{e}_k = \sum_{i \geq 2} \frac{c_i}{c_1} \left(\frac{\lambda_i}{\lambda_1} \right)^k \mathbf{v}_i.$$

Since $|\lambda_i/\lambda_1| \leq \rho$ for $i \geq 2$, the triangle inequality (Corollary 11.1.6) gives

$$\|\mathbf{e}_k\| \leq \rho^k \sum_{i \geq 2} \left| \frac{c_i}{c_1} \right| \|\mathbf{v}_i\| = C \rho^k \|\mathbf{v}_1\|.$$

Normalizing. Let $\omega_k = \overline{\sigma_k}$, where $\sigma_k = c_1 \lambda_1^k / |c_1 \lambda_1^k|$ has modulus 1. Then $\mathbf{A}^k \mathbf{x}^{(0)} = |c_1 \lambda_1^k| \sigma_k \mathbf{y}_k$, so $\mathbf{x}^{(k)} = \sigma_k \mathbf{y}_k / \|\mathbf{y}_k\|$ and $\omega_k \mathbf{x}^{(k)} = \mathbf{y}_k / \|\mathbf{y}_k\|$. For any nonzero $\mathbf{a}, \mathbf{b}$,

$$\begin{aligned} \frac{\mathbf{a}}{\|\mathbf{a}\|} - \frac{\mathbf{b}}{\|\mathbf{b}\|} &= \frac{\mathbf{a} - \mathbf{b}}{\|\mathbf{a}\|} + \mathbf{b} \left(\frac{1}{\|\mathbf{a}\|} - \frac{1}{\|\mathbf{b}\|} \right), \\ \left| \frac{1}{\|\mathbf{a}\|} - \frac{1}{\|\mathbf{b}\|} \right| \|\mathbf{b}\| &= \frac{|\|\mathbf{b}\| - \|\mathbf{a}\||}{\|\mathbf{a}\|} \leq \frac{\|\mathbf{a} - \mathbf{b}\|}{\|\mathbf{a}\|}, \end{aligned}$$

the last step by the reverse triangle inequality (Lemma 16.1.1). Hence $\|\mathbf{a}/\|\mathbf{a}\| - \mathbf{b}/\|\mathbf{b}\|\| \leq 2\|\mathbf{a} - \mathbf{b}\|/\|\mathbf{a}\|$. Apply this with $\mathbf{a} = \mathbf{y}_k$ and $\mathbf{b} = \mathbf{v}_1$. If $C\rho^k \leq \frac{1}{2}$ then $\|\mathbf{e}_k\| \leq \frac{1}{2}\|\mathbf{v}_1\|$, so $\|\mathbf{y}_k\| \geq \|\mathbf{v}_1\| - \|\mathbf{e}_k\| \geq \frac{1}{2}\|\mathbf{v}_1\|$, and

$$\|\omega_k \mathbf{x}^{(k)} - \mathbf{q}_1\| \leq \frac{2\|\mathbf{e}_k\|}{\|\mathbf{y}_k\|} \leq \frac{4\|\mathbf{e}_k\|}{\|\mathbf{v}_1\|} \leq 4C\rho^k,$$

which is (a).

The estimates. Write $\mathbf{z} = \omega_k \mathbf{x}^{(k)}$, a unit vector, and note $\mu_k = (\mathbf{x}^{(k)})^* \mathbf{A} \mathbf{x}^{(k)} = \mathbf{z}^* \mathbf{A} \mathbf{z}$ because $\overline{\omega_k} \omega_k = 1$. Since $\mathbf{A} \mathbf{q}_1 = \lambda_1 \mathbf{q}_1$ and $\|\mathbf{q}_1\| = 1$, we have $\mathbf{q}_1^* \mathbf{A} \mathbf{q}_1 = \lambda_1$, so

$$\mu_k - \lambda_1 = \mathbf{z}^* \mathbf{A} (\mathbf{z} - \mathbf{q}_1) + (\mathbf{z} - \mathbf{q}_1)^* \mathbf{A} \mathbf{q}_1.$$

By Cauchy–Schwarz (Theorem 11.1.5) and Theorem 16.3.2 (a), each term is at most $\|\mathbf{A}\| \|\mathbf{z} - \mathbf{q}_1\|$, so $|\mu_k - \lambda_1| \leq 2\|\mathbf{A}\| \|\mathbf{z} - \mathbf{q}_1\| \leq 8\|\mathbf{A}\| C \rho^k$ by (a). This proves the theorem. ■

Aftermath. One matrix–vector product per step: $2n^2$ flops for a dense $\mathbf{A}$, and far fewer when $\mathbf{A}$ is sparse — for the hyperlink matrix of Chapter 19 §08, one operation per link. Nothing is factorized and nothing is stored but a vector, which is why the method survives at sizes where the matrix itself could never be written down. What it buys with that cheapness is a *linear* rate: to gain one decimal digit takes $1/\log_{10}(1/\rho)$ steps, and for $\rho = 0.99$ that is 230 of them.

Example. Every iterate exactly

Let $\mathbf{A} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ and $\mathbf{x}^{(0)} = \mathbf{e}_1$. Compute $\mathbf{A}^k \mathbf{e}_1$, the estimates μ_k , and the errors, in closed form.

Solution. $\mathbf{A}$ has eigenvectors $(1, 1)$ and $(1, -1)$ with eigenvalues 3 and 1, so $\rho = \frac{1}{3}$. Writing $\mathbf{e}_1 = \frac{1}{2}(1, 1) + \frac{1}{2}(1, -1)$,

$$\mathbf{A}^k \mathbf{e}_1 = \frac{1}{2} 3^k (1, 1) + \frac{1}{2} (1, -1) = \left(\frac{3^k + 1}{2}, \frac{3^k - 1}{2} \right).$$

So the iterates are proportional to $(1, 0), (2, 1), (5, 4), (14, 13), (41, 40), (122, 121), \dots$, visibly running to the direction $(1, 1)$. Put $t = 3^k$ and $\mathbf{x} = (a, b)$ with $a = (t + 1)/2$, $b = (t - 1)/2$. Then $a^2 + b^2 = (t^2 + 1)/2$ and $ab = (t^2 - 1)/4$, so

$$\mu_k = \frac{2a^2 + 2ab + 2b^2}{a^2 + b^2} = \frac{(t^2 + 1) + \frac{1}{2}(t^2 - 1)}{\frac{1}{2}(t^2 + 1)} = \frac{3t^2 + 1}{t^2 + 1} = 3 - \frac{2}{9^k + 1}.$$

The values are 2, $\frac{14}{5}$, $\frac{122}{41}$, $\frac{1094}{365}$, $\frac{9842}{3281}$, with errors 1, $\frac{1}{5}$, $\frac{1}{41}$, $\frac{1}{365}$, $\frac{1}{3281}$ — falling by factors 5, 8.2, 8.9, 9.0, so by about a ninth in the limit, not a third.

The eigenvector error tells the other half. The component of $\mathbf{x} = (a, b)$ orthogonal to $(1, 1)$ is $\frac{1}{2}(a - b)(1, -1)$, of length $|a - b|/\sqrt{2} = 1/\sqrt{2}$, and $\|\mathbf{x}\| = \sqrt{(t^2 + 1)/2}$, so the sine of the angle between $\mathbf{x}^{(k)}$ and $(1, 1)$ is $(9^k + 1)^{-1/2}$, which does shrink like $3^{-k} = \rho^k$. Exactly:

$$3 - \mu_k = \frac{2}{9^k + 1} = 2 \sin^2 \theta_k.$$

The eigenvector converges at rate ρ^k and the eigenvalue estimate at rate ρ^{2k} , and the second is the square of the first. That is not an accident of this matrix.

The corollary explains it, and pays Chapter 17 §01 in its own words.

Corollary 24.4.2. Hermitian Matrices: the Rayleigh Quotient Converges Twice as Fast

Let $\mathbf{A} \in M_n(F)$ be Hermitian, with an orthonormal basis $\mathbf{q}_1, \dots, \mathbf{q}_n$ of eigenvectors, $\mathbf{A}\mathbf{q}_i = \lambda_i \mathbf{q}_i$, ordered by modulus so that $|\lambda_1| > |\lambda_2| \geq \dots \geq |\lambda_n|$. Let $\mathbf{x}^{(0)} = \sum_i c_i \mathbf{q}_i$ be a unit vector with $c_1 \neq 0$, and put $\rho = |\lambda_2/\lambda_1|$. Then for every k ,

$$|R_{\mathbf{A}}(\mathbf{x}^{(k)}) - \lambda_1| \leq D \rho^{2k}, \quad D = \frac{1}{|c_1|^2} \sum_{i \geq 2} |c_i|^2 |\lambda_i - \lambda_1|.$$

Idea. The Rayleigh quotient is unchanged by scaling, so evaluate it at the unnormalized $\mathbf{A}^k \mathbf{x}^{(0)}$, where the coefficients are just $c_i \lambda_i^k$. Both the numerator and the denominator are then sums of *nonnegative* multiples of λ_i^{2k} — the exponent is even, which is the whole point — so the denominator is bounded below by its first term alone, and the numerator, after λ_1 is subtracted, loses its first term entirely. What is left is λ_i^{2k} against λ_1^{2k} : one factor of ρ^2 per step.

Proof. Write $\mathbf{y} = \mathbf{A}^k \mathbf{x}^{(0)} = \sum_i c_i \lambda_i^k \mathbf{q}_i$. Since $c_1 \lambda_1^k \neq 0$, $\mathbf{y} \neq \mathbf{0}$, and $\mathbf{x}^{(k)}$ is a scalar multiple of $\mathbf{y}$ as in the proof of [Theorem 24.4.1](#), so $R_{\mathbf{A}}(\mathbf{x}^{(k)}) = R_{\mathbf{A}}(\mathbf{y})$ by the scale-invariance of the Rayleigh quotient (Definition 17.1.1). The basis is orthonormal and $\mathbf{A}\mathbf{y} = \sum_i c_i \lambda_i^{k+1} \mathbf{q}_i$, so by Theorem 11.2.1 (b) and (c),

$$R_{\mathbf{A}}(\mathbf{y}) = \frac{\langle \mathbf{A}\mathbf{y}, \mathbf{y} \rangle}{\langle \mathbf{y}, \mathbf{y} \rangle} = \frac{\sum_i |c_i|^2 \lambda_i^{2k+1}}{\sum_i |c_i|^2 \lambda_i^{2k}},$$

all the λ_i being real (Proposition 11.6.7). Subtract λ_1 , which is the same as subtracting λ_1 times the denominator from the numerator; the term $i = 1$ cancels, leaving

$$R_{\mathbf{A}}(\mathbf{y}) - \lambda_1 = \frac{\sum_{i \geq 2} |c_i|^2 \lambda_i^{2k} (\lambda_i - \lambda_1)}{\sum_i |c_i|^2 \lambda_i^{2k}}.$$

Every λ_i^{2k} is nonnegative, so the denominator is at least its first term $|c_1|^2 \lambda_1^{2k} = |c_1|^2 |\lambda_1|^{2k}$, and the numerator is at most $|\lambda_2|^{2k} \sum_{i \geq 2} |c_i|^2 |\lambda_i - \lambda_1|$ in modulus, since $|\lambda_i| \leq |\lambda_2|$ for $i \geq 2$. Dividing gives the stated bound with $\rho^{2k} = (|\lambda_2|/|\lambda_1|)^{2k}$. This proves the corollary. ■

This is the “second-order accuracy” of Chapter 17 §01, seen along an iteration: the eigenvector error is ρ^k and the eigenvalue error is its square. Theorem 20.6.2 says the same thing structurally and without any iteration — a residual of size ε yields an eigenvalue error of at most ε^2 divided by the gap — and the corollary above is that estimate made explicit for the particular unit vectors the power method produces.

Quick check. In [Example \(Every iterate exactly\)](#) the estimates μ_k increase towards 3 and never overshoot. Which earlier result says that they cannot?

Solution. Proposition 17.1.1: for a Hermitian $\mathbf{A}$, the values of $R_{\mathbf{A}}$ fill exactly the interval $[\lambda_n(\mathbf{A}), \lambda_1(\mathbf{A})]$ between the smallest and largest eigenvalues. Here $\mathbf{A} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ has eigenvalues 3 and 1, so every $\mu_k = R_{\mathbf{A}}(\mathbf{x}^{(k)})$ lies in $[1, 3]$ and no estimate can exceed 3. The power method on a Hermitian matrix with $\lambda_1 > 0$ dominant therefore produces a sequence of certified *lower* bounds for λ_1 , which is what Chapter 19 §04’s brackets were doing by hand.

Example. The method that never converges

Let $\mathbf{P} \in M_3(\mathbb{R})$ be the cyclic permutation matrix of Example (The Cyclic Permutation), with columns $\mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_1$, so that $\mathbf{P}(x_1, x_2, x_3) = (x_3, x_1, x_2)$. Run the power iteration from $\mathbf{x}^{(0)} = \mathbf{e}_1$ and from $\mathbf{x}^{(0)} = \frac{1}{\sqrt{14}}(1, 2, 3)$.

Solution. By Example (The Cyclic Permutation) the eigenvalues of $\mathbf{P}$ are the three cube roots of unity, all of modulus 1: the hypothesis $|\lambda_1| > |\lambda_2|$ fails as badly as it can. Every $\mathbf{P}^k$ is a permutation matrix, so no normalization is needed, and $\mathbf{P}^3 = \mathbf{I}$. From $\mathbf{e}_1$ the iterates are

$$\mathbf{e}_1 \rightarrow \mathbf{e}_2 \rightarrow \mathbf{e}_3 \rightarrow \mathbf{e}_1 \rightarrow \dots$$

and from $\frac{1}{\sqrt{14}}(1, 2, 3)$ they are $\frac{1}{\sqrt{14}}$ times $(1, 2, 3), (3, 1, 2), (2, 3, 1), (1, 2, 3), \dots$. In both cases the sequence takes three distinct values over and over, so neither it nor the line it spans has a limit. The estimates behave worse than that: they are *constant*, $\mu_k = 0$ for every k in the first case and $\mu_k = \frac{11}{14}$ in the second, so they converge perfectly well — to a number that is not an eigenvalue of $\mathbf{P}$ at all. Yet $\mathbf{P}$ does have the eigenvalue 1, with eigenvector $\mathbf{1}$, and the method never finds it.

The failure is exactly the one Chapter 19 §04 warned of for the swap $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, whose brackets “never improve”. It is not a defect of the arithmetic and no amount of care repairs it: the hypothesis of [Theorem 24.4.1](#) is false, and its conclusion is false too.

Warning. Two hypotheses, two different kinds of failure. If $|\lambda_1| = |\lambda_2|$, as in [Example \(The method that never converges\)](#), the method fails permanently and no starting vector helps. If instead $c_1 = 0$ — the starting vector happens to miss the dominant eigendirection — the method converges, but to the wrong eigenvector. The second failure is fragile and the first is not: a starting vector with $c_1 = 0$ lies in a proper subspace, and in the arithmetic of Definition 24.1.1 the first rounding error generally puts a small nonzero c_1 back, after which [Theorem 24.4.1](#) takes over, only C being large. This is the one situation in the chapter where rounding *helps*, and it is not a reason to rely on it: a starting vector chosen without regard to $\mathbf{A}$ is a better plan than a lucky rounding.

24.4.2 Shifting and inverting

The power method finds the eigenvalue of largest modulus, and that is usually not the one anybody wants. The repair is a change of matrix. If μ is not an eigenvalue of $\mathbf{A}$, then $\mathbf{A} - \mu\mathbf{I}$ is invertible, and the eigenvalues of its inverse are the numbers $(\lambda_i - \mu)^{-1}$ — **large exactly for the λ_i close to μ** . Multiplying repeatedly by $(\mathbf{A} - \mu\mathbf{I})^{-1}$ therefore singles out the eigenvalue nearest to μ , and μ is ours to choose.

Definition 24.4.2. Inverse Iteration

Let $\mathbf{A} \in M_n(F)$, let $\mu \in F$ be **not an eigenvalue of $\mathbf{A}$** , and let $\mathbf{x}^{(0)}$ be a unit vector. **Inverse iteration with shift μ** is the power iteration of [Definition 24.4.1](#) generated by $(\mathbf{A} - \mu\mathbf{I})^{-1}$ and $\mathbf{x}^{(0)}$; that is,

$$(\mathbf{A} - \mu\mathbf{I})\mathbf{z}^{(k)} = \mathbf{x}^{(k)}, \quad \mathbf{x}^{(k+1)} = \frac{\mathbf{z}^{(k)}}{\|\mathbf{z}^{(k)}\|}.$$

The eigenvalue estimates are taken from $\mathbf{A}$ itself: $\mu_k = (\mathbf{x}^{(k)})^* \mathbf{A} \mathbf{x}^{(k)}$.

The shifted matrix is **never inverted**. One factorization of $\mathbf{A} - \mu\mathbf{I}$ by elimination with pivoting is computed once, at a cost of $\frac{2}{3}n^3$ flops (Section 2), and each step then costs two triangular solves, $2n^2$ flops — the same price per step as the power method, after a one-off payment.

Theorem 24.4.3. Inverse Iteration Converges to the Nearest Eigenvector

Let $\mathbf{A} \in M_n(F)$ be diagonalizable with eigenbasis $\mathbf{v}_1, \dots, \mathbf{v}_n$ and eigenvalues $\lambda_1, \dots, \lambda_n$, and

let $\mu \in F$ be not an eigenvalue of $\mathbf{A}$. Suppose there is an index j with

$$|\lambda_j - \mu| < |\lambda_i - \mu| \quad \text{for every } i \neq j,$$

so that λ_j is the unique eigenvalue nearest μ . Let $\mathbf{x}^{(0)}$ be a unit vector whose $\mathbf{v}_j$ -coefficient is nonzero. Then inverse iteration with shift μ is defined for every k , the iterates approach the eigenline $\text{span}(\mathbf{v}_j)$ in the sense of [Theorem 24.4.1](#) (a), and $\mu_k \rightarrow \lambda_j$, both at the rate ρ_μ^k , where

$$\rho_\mu = \frac{|\lambda_j - \mu|}{\min_{i \neq j} |\lambda_i - \mu|} < 1.$$

Idea. There is nothing to prove that [Theorem 24.4.1](#) has not proved. The only work is to check that $(\mathbf{A} - \mu\mathbf{I})^{-1}$ satisfies its hypotheses — same eigenvectors, reciprocal shifted eigenvalues, a strict winner in modulus — and to read off what the ratio ρ of that theorem becomes after the reciprocals are taken: dividing by the largest reciprocal is multiplying by the smallest difference, so the ratio turns upside down.

Proof. Let $\mathbf{V} = (\mathbf{v}_1 \mid \cdots \mid \mathbf{v}_n)$, invertible because its columns are a basis, and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$, so that $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ ([Theorem 9.4.1](#)). Then

$$\mathbf{A} - \mu\mathbf{I} = \mathbf{V}(\mathbf{\Lambda} - \mu\mathbf{I})\mathbf{V}^{-1},$$

and $\mathbf{\Lambda} - \mu\mathbf{I}$ is invertible because no λ_i equals μ . Hence $\mathbf{B} := (\mathbf{A} - \mu\mathbf{I})^{-1} = \mathbf{V}(\mathbf{\Lambda} - \mu\mathbf{I})^{-1}\mathbf{V}^{-1}$ is diagonalizable, with the **same** eigenvectors $\mathbf{v}_i$ and eigenvalues $\beta_i = (\lambda_i - \mu)^{-1}$.

Their moduli are $|\beta_i| = 1/|\lambda_i - \mu|$, so the hypothesis on j says exactly that $|\beta_j| > |\beta_i|$ for every $i \neq j$: the eigenvalue β_j is strictly dominant. Relabeling so that j comes first, the second largest modulus is $\max_{i \neq j} |\beta_i| = 1/\min_{i \neq j} |\lambda_i - \mu|$, so the ratio of [Theorem 24.4.1](#) is

$$\frac{\max_{i \neq j} |\beta_i|}{|\beta_j|} = \frac{|\lambda_j - \mu|}{\min_{i \neq j} |\lambda_i - \mu|} = \rho_\mu.$$

Since the $\mathbf{v}_j$ -coefficient of $\mathbf{x}^{(0)}$ is nonzero, [Theorem 24.4.1](#) applies to $\mathbf{B}$ and $\mathbf{x}^{(0)}$ and gives unimodular ω_k with $\|\omega_k \mathbf{x}^{(k)} - \mathbf{q}_j\| \leq 4C\rho_\mu^k$, where $\mathbf{q}_j = \mathbf{v}_j/\|\mathbf{v}_j\|$ and C is that theorem's constant for $\mathbf{B}$ and $\mathbf{x}^{(0)}$. That is the convergence of the line.

For the estimates, repeat the last paragraph of the proof of [Theorem 24.4.1](#) with $\mathbf{A}$ in place of $\mathbf{B}$: with $\mathbf{z} = \omega_k \mathbf{x}^{(k)}$ we have $\mu_k = \mathbf{z}^* \mathbf{A} \mathbf{z}$ and $\mathbf{q}_j^* \mathbf{A} \mathbf{q}_j = \lambda_j$, so $|\mu_k - \lambda_j| \leq 2\|\mathbf{A}\|\|\mathbf{z} - \mathbf{q}_j\| \leq 8\|\mathbf{A}\|C\rho_\mu^k$. This proves the theorem. ■

Example. Aiming at the smaller eigenvalue

Let $\mathbf{A} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ again, and take the shift $\mu = \frac{3}{4}$, which is nearer to the eigenvalue 1 than to 3. Run inverse iteration from $\mathbf{x}^{(0)} = \mathbf{e}_1$, and compare with [Example \(Every iterate exactly\)](#).

Solution. $\mathbf{A} - \frac{3}{4}\mathbf{I} = \begin{pmatrix} 5/4 & 1 \\ 1 & 5/4 \end{pmatrix}$ has determinant $\frac{25}{16} - 1 = \frac{9}{16}$, so

$$(\mathbf{A} - \frac{3}{4}\mathbf{I})^{-1} = \frac{16}{9} \begin{pmatrix} 5/4 & -1 \\ -1 & 5/4 \end{pmatrix} = \frac{1}{9} \begin{pmatrix} 20 & -16 \\ -16 & 20 \end{pmatrix}.$$

Ignoring positive scalars, each step is $(a, b) \mapsto (5a - 4b, -4a + 5b)$, so the iterates are

$$(1, 0) \rightarrow (5, -4) \rightarrow (41, -40) \rightarrow (365, -364) \rightarrow (3281, -3280) \rightarrow \dots,$$

running to the direction $(1, -1)$ — the eigenvector of the eigenvalue 1. In closed form: the eigenvalues of $(\mathbf{A} - \frac{3}{4}\mathbf{I})^{-1}$ are $(3 - \frac{3}{4})^{-1} = \frac{4}{9}$ and $(1 - \frac{3}{4})^{-1} = 4$, so expanding $\mathbf{e}_1 = \frac{1}{2}(1, 1) + \frac{1}{2}(1, -1)$ gives

$$(\mathbf{A} - \frac{3}{4}\mathbf{I})^{-k} \mathbf{e}_1 = \frac{1}{2}(\frac{4}{9})^k(1, 1) + \frac{1}{2}4^k(1, -1) \propto (9^k + 1, -(9^k - 1)),$$

the last step on multiplying by $2(9/4)^k$. Halving, this is the list above. Writing $\mathbf{x} = (a, b)$ with $a + b = 1$, we have $2ab = 1 - (a^2 + b^2)$ and so

$$\mu_k = \frac{2a^2 + 2ab + 2b^2}{a^2 + b^2} = 1 + \frac{1}{a^2 + b^2} = 1 + \frac{2}{81^k + 1},$$

since $a^2 + b^2 = \frac{1}{2}(81^k + 1)$. The errors are $1, \frac{1}{41}, \frac{1}{3281}, \frac{1}{265721}, \frac{1}{21523361}, \dots$, each about $\frac{1}{81}$ of the last, while the sine of the angle to $(1, -1)$ is $(81^k + 1)^{-1/2}$, shrinking by $\frac{1}{9}$ per step. The theorem predicts exactly that:

$$\rho_\mu = \frac{|1 - \frac{3}{4}|}{|3 - \frac{3}{4}|} = \frac{1/4}{9/4} = \frac{1}{9},$$

with the estimates again going at the squared rate $\rho_\mu^{2k} = 81^{-k}$, as in [Example \(Every iterate exactly\)](#) and for the same reason. Compare that example, where the same matrix and the same starting vector gave rate $\frac{1}{3}$ towards the *other* eigenvector: one well-chosen shift has both redirected the method and made it gain twice as many digits per step.

The objection, and the answer. If μ is close to λ_j , then $\mathbf{A} - \mu\mathbf{I}$ is close to singular and $\kappa_2(\mathbf{A} - \mu\mathbf{I})$ is enormous. For normal $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{U}^*$ (Corollary 12.4.3) we have $\mathbf{A} - \mu\mathbf{I} = \mathbf{U}(\mathbf{D} - \mu\mathbf{I})\mathbf{U}^*$, whose singular values are the numbers $|\lambda_i - \mu|$, so by Proposition 16.8.2 (d)

$$\kappa_2(\mathbf{A} - \mu\mathbf{I}) = \frac{\max_i |\lambda_i - \mu|}{|\lambda_j - \mu|},$$

which grows without bound as $\mu \rightarrow \lambda_j$. Section 3 spent its length on how much a large condition number costs. And the better the shift, the worse the conditioning: **the method is deliberately solving an ill-conditioned system, and it does so on purpose.**

It does no harm, and here is why. Suppose the linear solve is backward stable in the sense of Definition 24.1.3, so that the computed $\hat{\mathbf{z}}$ is the exact solution of

$$(\mathbf{A} - \mu\mathbf{I} + \mathbf{\Delta})\hat{\mathbf{z}} = \mathbf{x}, \quad \|\mathbf{\Delta}\| \leq cu \|\mathbf{A} - \mu\mathbf{I}\|,$$

for a unit $\mathbf{x}$. The computed $\hat{\mathbf{z}}$ may be wrong in every digit — the forward error is bounded only by $\kappa_2(\mathbf{A} - \mu\mathbf{I})cu$, which may exceed 1. But rearrange the same equation as a residual: $(\mathbf{A} - \mu\mathbf{I})\hat{\mathbf{z}} - \mathbf{x} = -\mathbf{\Delta}\hat{\mathbf{z}}$, so the **normalized** vector $\hat{\mathbf{q}} = \hat{\mathbf{z}}/\|\hat{\mathbf{z}}\|$ satisfies

$$\|\mathbf{A}\hat{\mathbf{q}} - \mu\hat{\mathbf{q}}\| \leq \frac{1}{\|\hat{\mathbf{z}}\|} + cu \|\mathbf{A} - \mu\mathbf{I}\|.$$

The first term is small when $\|\hat{\mathbf{z}}\|$ is large, and that takes an argument and a hypothesis. Near-singularity alone will not supply it: for $\mathbf{A} = \text{diag}(1, 2)$, $\mu = 1 + 10^{-10}$, and $\mathbf{x} = \mathbf{e}_2$, the matrix $\mathbf{A} - \mu\mathbf{I}$ is as near singular as one pleases and yet $\mathbf{z} = (\mathbf{A} - \mu\mathbf{I})^{-1}\mathbf{x}$ has norm 1 to ten digits. What is needed is that $\mathbf{x}$ see the near-singular direction.

Take $\mathbf{A}$ normal with orthonormal eigenbasis $\mathbf{q}_1, \dots, \mathbf{q}_n$ (Corollary 12.4.3), so that $\mathbf{A}^*\mathbf{q}_j = \bar{\lambda}_j\mathbf{q}_j$ and hence $\mathbf{q}_j^*(\mathbf{A} - \mu\mathbf{I}) = (\lambda_j - \mu)\mathbf{q}_j^*$. Multiplying the perturbed equation by $\mathbf{q}_j^*$ gives $\mathbf{q}_j^*\mathbf{x} = (\lambda_j - \mu)\mathbf{q}_j^*\hat{\mathbf{z}} + \mathbf{q}_j^*\Delta\hat{\mathbf{z}}$, so by Cauchy–Schwarz (Theorem 11.1.5) and Theorem 16.3.2 (a),

$$\|\hat{\mathbf{z}}\| \geq \frac{|\mathbf{q}_j^*\mathbf{x}|}{|\lambda_j - \mu| + \|\Delta\|}.$$

So a shift close to λ_j does force $\|\hat{\mathbf{z}}\|$ to be large, **provided the component $\mathbf{q}_j^*\mathbf{x}$ is not small** — which is the hypothesis [Theorem 24.4.3](#) already makes on the starting vector. Granting it, the first term of the residual bound is small, and the second is at the rounding level. **The residual of the normalized answer is small even though the answer itself is inaccurate**, and a small residual is all that an eigenvector approximation has to have.

For Hermitian $\mathbf{A}$ one can say precisely how much it buys. Let $\hat{\mathbf{q}}$ be a unit vector, $\hat{\mu} = R_{\mathbf{A}}(\hat{\mathbf{q}})$ its Rayleigh quotient, and $\varepsilon = \|\mathbf{A}\hat{\mathbf{q}} - \hat{\mu}\hat{\mathbf{q}}\|$ its residual. Expand $\hat{\mathbf{q}} = \sum_i d_i\mathbf{q}_i$ in an orthonormal eigenbasis (Corollary 12.4.3); by Theorem 11.2.1 (c),

$$\varepsilon^2 = \sum_i |d_i|^2 (\lambda_i - \hat{\mu})^2.$$

If λ_j is the only eigenvalue within distance g of $\hat{\mu}$, every term with $i \neq j$ has $(\lambda_i - \hat{\mu})^2 \geq g^2$, so $\sum_{i \neq j} |d_i|^2 \leq \varepsilon^2/g^2$ and the sine of the angle between $\hat{\mathbf{q}}$ and $\mathbf{q}_j$ is at most ε/g . Together with Theorem 20.6.1, which puts an eigenvalue of $\mathbf{A}$ within ε of $\hat{\mu}$, and Theorem 20.6.2, which improves that to ε^2/g , the small residual certifies both halves of the answer.

Warning. “**Ill-conditioned**” is a statement about the wrong question. The system $(\mathbf{A} - \mu\mathbf{I})\mathbf{z} = \mathbf{x}$ is indeed almost unsolvable *if what is wanted is $\mathbf{z}$* . What is wanted is the direction of $\mathbf{z}$, and the direction is precisely the part the ill-conditioning leaves alone: the huge factor $1/(\lambda_j - \mu)$ multiplies the component along $\mathbf{v}_j$, which is the component being sought, and dividing by $\|\hat{\mathbf{z}}\|$ removes it. The residual estimate above is the rigorous form of that sentence. The same reasoning does **not** license inverse iteration when the eigenvalue is not simple or two eigenvalues are nearly equal: then g is small, the bound ε/g says nothing, and the computed direction really is unreliable.

24.4.3 The shift that improves itself

Inverse iteration needs a shift, and a better shift converges faster. But the iteration is manufacturing better and better approximations to λ_j as it goes, in the estimates μ_k . Use them.

Definition 24.4.3. Rayleigh Quotient Iteration

Let $\mathbf{A} \in M_n(F)$ be Hermitian and let $\mathbf{x}^{(0)}$ be a unit vector. **Rayleigh quotient iteration** is the sequence defined for $k \geq 0$ by

$$\mu_k = R_{\mathbf{A}}(\mathbf{x}^{(k)}) = (\mathbf{x}^{(k)})^* \mathbf{A} \mathbf{x}^{(k)}, \quad (\mathbf{A} - \mu_k \mathbf{I}) \mathbf{z}^{(k)} = \mathbf{x}^{(k)}, \quad \mathbf{x}^{(k+1)} = \frac{\mathbf{z}^{(k)}}{\|\mathbf{z}^{(k)}\|},$$

the step being **undefined, and the iteration stopped, when μ_k is an eigenvalue of $\mathbf{A}$** — in which case μ_k is the answer, exactly.

It is inverse iteration with the shift refreshed at every step, and it is what Chapter 17 §01 and Chapter 20 §06 both promised. The price is a new factorization of $\mathbf{A} - \mu_k \mathbf{I}$ at every step, $\frac{2}{3}n^3$ flops instead of $2n^2$. What is bought is the following rate.

Theorem 24.4.4. Rayleigh Quotient Iteration Converges Cubically

Let $\mathbf{A} \in M_n(F)$ be Hermitian with an orthonormal basis $\mathbf{q}_1, \dots, \mathbf{q}_n$ of eigenvectors, $\mathbf{A}\mathbf{q}_i = \lambda_i \mathbf{q}_i$, and fix an index j whose eigenvalue is **simple**. Put

$$g = \min_{i \neq j} |\lambda_i - \lambda_j| > 0, \quad D = \max_{i \neq j} |\lambda_i - \lambda_j|,$$

and for a unit vector $\mathbf{x}$ write $s(\mathbf{x}) = \|\mathbf{x} - \mathbf{q}_j \mathbf{q}_j^* \mathbf{x}\|$, the length of the part of $\mathbf{x}$ orthogonal to $\mathbf{q}_j$.

Let $\mathbf{x} = \mathbf{x}^{(k)}$ be a unit iterate with $s = s(\mathbf{x})$ satisfying $s \leq \frac{1}{2}$ and $Ds^2 \leq \frac{g}{2}$. Then $|\mu_k - \lambda_j| \leq Ds^2$, and either $\mu_k = \lambda_j$ — the iteration has found the eigenvalue exactly — or the next iterate is defined and

$$s(\mathbf{x}^{(k+1)}) \leq \frac{4D}{g} s^3.$$

Consequently, if $s_0 = s(\mathbf{x}^{(0)})$ satisfies $s_0 \leq \frac{1}{2}$ and $s_0^2 \leq g/(8D)$, then $a_k := \sqrt{4D/g} s(\mathbf{x}^{(k)})$ satisfies $a_k \leq a_0^3$ with $a_0 \leq 1/\sqrt{2}$, so $s(\mathbf{x}^{(k)}) \rightarrow 0$ and $\mu_k \rightarrow \lambda_j$.

Idea. Two facts, multiplied together. The Rayleigh quotient is accurate to *second* order in the eigenvector error, so a shift built from an iterate at distance s is wrong by only Ds^2 — that is [Corollary 24.4.2](#)'s squaring again, in its noniterative form. And one step of inverse iteration with a shift wrong by δ multiplies the error by about δ/g . Put $\delta = Ds^2$ and the step multiplies s by Ds^2/g : a factor s^2 per step, which is one cube. Everything else is the arithmetic of keeping the shift far enough from the other eigenvalues for g to survive, and that is what the two smallness hypotheses are for.

Proof. Expand $\mathbf{x} = \sum_i c_i \mathbf{q}_i$ with $\sum_i |c_i|^2 = 1$ (Theorem 11.2.1 (c)). Then $\mathbf{q}_j \mathbf{q}_j^* \mathbf{x} = c_j \mathbf{q}_j$, so

$$s^2 = \sum_{i \neq j} |c_i|^2, \quad |c_j|^2 = 1 - s^2.$$

Step 1: the shift is accurate to second order. By Theorem 11.2.1 (b), $\mu_k = \mathbf{x}^* \mathbf{A} \mathbf{x} = \sum_i |c_i|^2 \lambda_i$, and since the weights sum to 1,

$$\mu_k - \lambda_j = \sum_{i \neq j} |c_i|^2 (\lambda_i - \lambda_j), \quad |\mu_k - \lambda_j| \leq D \sum_{i \neq j} |c_i|^2 = Ds^2.$$

Step 2: the shift stays away from the other eigenvalues. By hypothesis $Ds^2 \leq g/2$, so $|\mu_k - \lambda_j| \leq g/2$ and, for $i \neq j$,

$$|\lambda_i - \mu_k| \geq |\lambda_i - \lambda_j| - |\lambda_j - \mu_k| \geq g - \frac{g}{2} = \frac{g}{2} > 0.$$

In particular the only eigenvalue μ_k can equal is λ_j . If it does, the iteration stops with the exact answer. Otherwise $\mathbf{A} - \mu_k \mathbf{I}$ is invertible and the step is defined.

Step 3: one step. Since $\mathbf{A}\mathbf{q}_i = \lambda_i \mathbf{q}_i$,

$$\mathbf{z}^{(k)} = (\mathbf{A} - \mu_k \mathbf{I})^{-1} \mathbf{x} = \sum_i \frac{c_i}{\lambda_i - \mu_k} \mathbf{q}_i.$$

The quantity $s(\cdot)$ is unchanged by scaling, so $s(\mathbf{x}^{(k+1)}) = s(\mathbf{z}^{(k)} / \|\mathbf{z}^{(k)}\|)$ is the ratio of the length of the part of $\mathbf{z}^{(k)}$ orthogonal to $\mathbf{q}_j$ to the length of $\mathbf{z}^{(k)}$:

$$s(\mathbf{x}^{(k+1)})^2 = \frac{\sum_{i \neq j} \frac{|c_i|^2}{|\lambda_i - \mu_k|^2}}{\sum_i \frac{|c_i|^2}{|\lambda_i - \mu_k|^2}}.$$

Bound the numerator above using Step 2, and the denominator below by its $i = j$ term alone:

$$s(\mathbf{x}^{(k+1)})^2 \leq \frac{(2/g)^2 s^2}{|c_j|^2 / |\lambda_j - \mu_k|^2} = \frac{4s^2 |\lambda_j - \mu_k|^2}{g^2 (1 - s^2)} \leq \frac{4s^2 D^2 s^4}{g^2 (1 - s^2)},$$

using Step 1 at the end. Taking square roots and using $s \leq \frac{1}{2}$, so that $\sqrt{1 - s^2} \geq \frac{\sqrt{3}}{2} \geq \frac{1}{2}$,

$$s(\mathbf{x}^{(k+1)}) \leq \frac{2D s^3}{g \sqrt{1 - s^2}} \leq \frac{4D}{g} s^3.$$

Step 4: the iteration. Put $C = 4D/g$ and $a_k = \sqrt{C} s(\mathbf{x}^{(k)})$. The hypotheses $s_0 \leq \frac{1}{2}$ and $s_0^2 \leq g/(8D)$ give $a_0^2 = C s_0^2 \leq \frac{1}{2}$, so $a_0 \leq 1/\sqrt{2} < 1$, and also $D s_0^2 \leq g/2$. Whenever the hypotheses hold at step k , Step 3 gives $a_{k+1} \leq \sqrt{C} C s_k^3 = a_k^3$, so $a_{k+1} \leq a_k^3 \leq a_k \leq a_0$; hence $s_{k+1} \leq s_k \leq s_0 \leq \frac{1}{2}$ and $D s_{k+1}^2 \leq D s_0^2 \leq g/2$, and the hypotheses hold again. By induction $a_k \leq a_0^{3^k}$ for every k . Since $a_0 < 1$, $a_k \rightarrow 0$ and so $s(\mathbf{x}^{(k)}) \rightarrow 0$; by Step 1, $|\mu_k - \lambda_j| \leq D s_k^2 \rightarrow 0$. This proves the theorem. ■

Cubic convergence is a different animal from the linear convergence of [Theorem 24.4.1](#). Each step roughly **triples the number of correct digits**, so an iterate with two correct digits has six after one step and eighteen after two — which is more than the arithmetic holds. In practice the iteration is over in three or four steps from any reasonable start, and the cost is dominated by those three or four factorizations.

Example. Three steps, exactly

Run Rayleigh quotient iteration on $\mathbf{A} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ from $\mathbf{x}^{(0)} \propto (2, 1)$, in exact arithmetic.

Solution. Write each iterate as an integer vector, scaling being irrelevant. Take $\mathbf{q}_j = \frac{1}{\sqrt{2}}(1, 1)$, so $\lambda_j = 3$, and here $g = D = 2$, giving $4D/g = 4$.

Step 0. $\mathbf{x} \propto (2, 1)$. Then $\mu_0 = (2 \cdot 2^2 + 2 \cdot 2 \cdot 1 + 2 \cdot 1^2)/5 = \frac{14}{5}$, and $\mathbf{A} - \frac{14}{5} \mathbf{I} = \begin{pmatrix} -4/5 & 1 \\ 1 & -4/5 \end{pmatrix}$ has determinant $-\frac{9}{25}$. Solving gives $\mathbf{z} \propto (13, 14)$.

Step 1. $\mathbf{x} \propto (13, 14)$, $\mu_1 = \frac{1094}{365}$, and $3 - \mu_1 = \frac{1}{365}$. Solving again gives $\mathbf{z} \propto (9842, 9841)$.

Step 2. $\mu_2 = \frac{581130734}{193710245}$, with $3 - \mu_2 = \frac{1}{193710245}$.

Step 3. $\mathbf{x} \propto (3812798742493, 3812798742494)$ and $3 - \mu_3 = \frac{1}{29074868501520029845195085}$.

The eigenvalue errors are

$$2 \cdot 10^{-1}, \quad 2.7 \cdot 10^{-3}, \quad 5.2 \cdot 10^{-9}, \quad 3.4 \cdot 10^{-26},$$

and the eigenvector errors s_k are

$$3.16 \cdot 10^{-1}, \quad 3.70 \cdot 10^{-2}, \quad 5.08 \cdot 10^{-5}, \quad 1.31 \cdot 10^{-13}.$$

Check the theorem: $s_1/s_0^3 = 1.17$, $s_2/s_1^3 = 1.00$, $s_3/s_2^3 = 1.00$, all comfortably below the predicted constant $4D/g = 4$. Four steps have produced twenty-five correct digits from a starting vector accurate to none.

Remark. **The non-Hermitian rate is quadratic, and we do not prove it.** For a general diagonalizable $\mathbf{A}$, with μ_k taken to be $(\mathbf{x}^{(k)})^* \mathbf{A} \mathbf{x}^{(k)}$, the same iteration converges locally at a *quadratic* rate rather than a cubic one. The reason is visible in Step 1 above: it used that μ_k is a weighted average of the eigenvalues with weights $|c_i|^2$, so that the first-order term in s is absent — and that identity is exactly the Hermitian spectral theorem at work. Without it the shift is accurate only to first order, $|\mu_k - \lambda_j| = O(s)$, and the step multiplies s by $O(s)$. **The general statement is classical and is not proved here**, and nothing in this book depends on it.

24.4.4 Deflation, and what comes next

One eigenpair is rarely enough. For Hermitian $\mathbf{A}$, once a unit eigenvector $\mathbf{q}_1$ and its eigenvalue λ_1 are known, the next one can be reached by removing the first: put

$$\mathbf{A}' = \mathbf{A} - \lambda_1 \mathbf{q}_1 \mathbf{q}_1^*.$$

Let $\mathbf{q}_1, \dots, \mathbf{q}_n$ be an orthonormal eigenbasis (Corollary 12.4.3). Then $\mathbf{A}' \mathbf{q}_1 = \lambda_1 \mathbf{q}_1 - \lambda_1 \mathbf{q}_1 (\mathbf{q}_1^* \mathbf{q}_1) = \mathbf{0}$, and for $i \neq 1$, $\mathbf{A}' \mathbf{q}_i = \lambda_i \mathbf{q}_i - \lambda_1 \mathbf{q}_1 (\mathbf{q}_1^* \mathbf{q}_i) = \lambda_i \mathbf{q}_i$ by orthonormality. So $\mathbf{A}'$ is Hermitian with the same eigenvectors and the eigenvalue list $0, \lambda_2, \dots, \lambda_n$: the power method applied to $\mathbf{A}'$ finds λ_2 , provided λ_2 is now dominant in modulus.

Warning. **Deflation is exact only with an exact eigenvector.** If the computed $\hat{\mathbf{q}}_1$ differs from $\mathbf{q}_1$ by an angle η , then the matrix actually formed differs from the intended $\mathbf{A}'$ by $\lambda_1 (\mathbf{q}_1 \mathbf{q}_1^* - \hat{\mathbf{q}}_1 \hat{\mathbf{q}}_1^*)$, of norm of order $|\lambda_1| \eta$, and by the perturbation bounds of Chapters 17 and 20 that error is inherited by every eigenvalue and eigenvector computed afterwards. Deflating $n - 1$ times accumulates $n - 1$ such errors. For one or two extreme eigenvalues deflation is the right tool; for the whole spectrum it is not, and the next section builds the algorithm that is.

Everything in this section has been a single vector chasing a single eigenvector. The next section

iterates on a whole basis at once, and the result — the QR algorithm — turns out to be the power method in disguise, its accumulated factors furnishing a **QR** factorization of $\mathbf{A}^k$, with the shifts of this section reappearing as the device that makes it fast.

24.4.5 Exercises

24.4.5.1 A. Check your understanding

Exercise. A1

- State the two hypotheses of [Theorem 24.4.1](#) on the eigenvalues and on the starting vector, and say what the rate ρ is.
- For a Hermitian $\mathbf{A}$, at what rate does the eigenvector converge and at what rate does the Rayleigh quotient? Which earlier result explains the relation between the two?
- Explain in one sentence why inverse iteration with shift μ finds the eigenvalue nearest μ .
- True or false: “the shifted matrix in inverse iteration is nearly singular, so the method is unreliable”. Justify your answer.
- Give the per-step cost of the power method, of inverse iteration, and of Rayleigh quotient iteration, for a dense $n \times n$ matrix.

Solution.

- The dominant eigenvalue must be strictly dominant **in modulus**, $|\lambda_1| > |\lambda_2|$, and the starting vector must have a nonzero coefficient c_1 on the corresponding eigenvector. The rate is $\rho = |\lambda_2/\lambda_1|$, the ratio of the two largest moduli.
- The eigenvector converges at rate ρ^k ([Theorem 24.4.1](#) (a)) and the Rayleigh quotient at rate ρ^{2k} ([Corollary 24.4.2](#)). The relation is the second-order accuracy of the Rayleigh quotient at a Hermitian matrix, quantified by Theorem 20.6.2: a residual of size ε gives an eigenvalue error of at most ε^2 divided by the gap.
- Because $(\mathbf{A} - \mu\mathbf{I})^{-1}$ has the same eigenvectors as $\mathbf{A}$ and the eigenvalues $(\lambda_i - \mu)^{-1}$, whose moduli are largest precisely for the λ_i nearest μ ([Theorem 24.4.3](#)).
- False. The forward error in $\mathbf{z}^{(k)}$ is indeed large, but only its **direction** is used, and normalizing removes the part of the error that the ill-conditioning creates; the residual $\|\mathbf{A}\hat{\mathbf{q}} - \mu\hat{\mathbf{q}}\|$ of the normalized answer is small, and that is what certifies an eigenvector. The method is unreliable for a different reason — a small gap g — as the warning after [Theorem 24.4.3](#) says.
- Power method: one matrix–vector product, $2n^2$ flops. Inverse iteration: one factorization of $\mathbf{A} - \mu\mathbf{I}$ by elimination with pivoting, $\frac{2}{3}n^3$, computed once, then $2n^2$ per step. Rayleigh quotient iteration: a new factorization every step, $\frac{2}{3}n^3$ per step — paid for by the cubic rate, which needs only three or four steps.

24.4.5.2 B. Practice

Exercise. B1: A rate, read off the spectrum

Let $\mathbf{A} = \begin{pmatrix} 4 & 1 \\ 2 & 3 \end{pmatrix}$.

- (a) Find the eigenvalues and eigenvectors of $\mathbf{A}$, and the rate ρ of the power method.
- (b) Starting from $\mathbf{x}^{(0)} = \mathbf{e}_1$, compute $\mathbf{A}^k \mathbf{e}_1$ in closed form and verify that the direction converges at rate ρ^k .
- (c) How many steps are needed to gain six decimal digits?

Solution.

- (a) $\text{tr } \mathbf{A} = 7$ and $\det \mathbf{A} = 10$, so the eigenvalues are 5 and 2. For $\lambda = 5$: $(\mathbf{A} - 5\mathbf{I})\mathbf{v} = \mathbf{0}$ reads $-v_1 + v_2 = 0$, giving $\mathbf{v}_1 = (1, 1)$. For $\lambda = 2$: $2v_1 + v_2 = 0$, giving $\mathbf{v}_2 = (1, -2)$. Hence $\rho = \frac{2}{5}$.
- (b) $\mathbf{e}_1 = \frac{2}{3}(1, 1) + \frac{1}{3}(1, -2)$, as one checks by adding. So

$$\mathbf{A}^k \mathbf{e}_1 = \frac{2}{3}5^k(1, 1) + \frac{1}{3}2^k(1, -2).$$

Dividing by $\frac{2}{3}5^k$ gives $(1, 1) + \frac{1}{2}(\frac{2}{5})^k(1, -2)$, so the distance from the direction $(1, 1)$ is a constant times $(\frac{2}{5})^k = \rho^k$, as [Theorem 24.4.1](#) says.

- (c) Each step multiplies the error by $\rho = 0.4$, that is by $10^{-\log_{10}(2.5)} = 10^{-0.398}$. Six digits need $6/0.398 = 15.1$ steps, so sixteen (and indeed $\rho^{15} = 1.1 \cdot 10^{-6}$, $\rho^{16} = 4.3 \cdot 10^{-7}$).

Exercise. B2: Choosing a shift

Let $\mathbf{A} = \text{diag}(10, 3, 1)$.

- (a) What is the rate of the plain power method?
- (b) Find the rate of inverse iteration with shift $\mu = 0$, and say which eigenvector it finds.
- (c) Find the rate of inverse iteration with shift $\mu = \frac{5}{2}$, and say which eigenvector it finds.
- (d) Which shift would you choose to compute the middle eigenvalue as fast as possible, and what rate does it give? Explain why no shift gives a rate of 0.

Solution.

- (a) By [Theorem 24.4.1](#), $\rho = |3/10| = 0.3$, and the iterates converge to $\mathbf{e}_1$.
- (b) With $\mu = 0$ the eigenvalue nearest 0 is 1, at distance 1; the next nearest is 3, at distance 3. So $\rho_\mu = \frac{1}{3}$, and the method finds $\mathbf{e}_3$. (Inverse iteration with shift 0 is the power method on $\mathbf{A}^{-1}$.)
- (c) Distances from $\frac{5}{2}$ are 7.5, 0.5, 1.5. The nearest eigenvalue is 3, the next nearest 1, so $\rho_\mu = 0.5/1.5 = \frac{1}{3}$, and the method finds $\mathbf{e}_2$, the **middle** eigenvector — which the plain power method can never reach.
- (d) Take $\mu = 3 + \delta$ with $0 < \delta < 1$: the distances are $7 - \delta$, δ , $2 + \delta$, so $\rho_\mu = \delta/(2 + \delta) \rightarrow 0$ as $\delta \rightarrow 0$. Any μ very close to 3 is excellent. No shift gives rate exactly 0, because $\mu = 3$ is an eigenvalue and [Definition 24.4.2](#) then does not apply: $\mathbf{A} - 3\mathbf{I}$ is singular and there is nothing to solve with. The limiting case is not a fast method but no method — which is why Rayleigh quotient iteration, which walks the shift towards the eigenvalue, stops the moment it arrives.

Exercise. B3: Deflating a symmetric matrix

Let $\mathbf{A} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$.

- Write down $\mathbf{A}' = \mathbf{A} - 3\mathbf{q}_1\mathbf{q}_1^\top$ for the unit eigenvector $\mathbf{q}_1$ of the eigenvalue 3, and verify its eigenvalues directly.
- What does the power method of [Theorem 24.4.1](#), applied to $\mathbf{A}'$ from $\mathbf{x}^{(0)} = \mathbf{e}_1$, produce, and at what rate?
- Explain why the answer to (b) makes deflation look better than it is.

Solution.

- $\mathbf{q}_1 = \frac{1}{\sqrt{2}}(1, 1)$ and $\mathbf{q}_1\mathbf{q}_1^\top = \frac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$, so

$$\mathbf{A}' = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} - \frac{3}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

Then $\mathbf{A}'(1, 1) = \mathbf{0}$ and $\mathbf{A}'(1, -1) = (1, -1)$: the eigenvalues are 0 and 1, as predicted, with the same eigenvectors.

- $\mathbf{A}'\mathbf{e}_1 = \frac{1}{2}(1, -1)$, which is already an eigenvector, so the power iteration of [Definition 24.4.1](#) gives $\mathbf{x}^{(1)} = \frac{1}{\sqrt{2}}(1, -1)$ and every later iterate equals it: it converges in one step, with $\mu_1 = 1$ exactly.
- Because the deflation was exact. The eigenvector $\mathbf{q}_1 = \frac{1}{\sqrt{2}}(1, 1)$ was known in closed form, so $\mathbf{A}'$ had the eigenvalue 0 exactly and the remaining eigenvector survived untouched. In a computation $\mathbf{q}_1$ is known only to within the accuracy of the first run, and by the warning after the deflation formula the error in $\mathbf{q}_1$ enters $\mathbf{A}'$ as a perturbation of size about $|\lambda_1|\eta$, which every later eigenpair inherits. One clean step here is not evidence about $n - 1$ dirty ones.

24.4.5.3 C. Going deeper**Exercise. C1: What the squaring really needs**

Let $\mathbf{A} \in M_n(F)$ be Hermitian with orthonormal eigenbasis $\mathbf{q}_1, \dots, \mathbf{q}_n$, and let $\mathbf{x}$ be a unit vector with $s = \|\mathbf{x} - \mathbf{q}_1\mathbf{q}_1^*\mathbf{x}\|$.

- Prove that $|R_{\mathbf{A}}(\mathbf{x}) - \lambda_1| \leq s^2 \max_{i \neq 1} |\lambda_i - \lambda_1|$.
- Show by example that the corresponding statement for a non-Hermitian $\mathbf{A}$ is false: exhibit $\mathbf{A} \in M_2(\mathbb{C})$, an eigenvalue λ_1 with unit eigenvector $\mathbf{q}_1$, and unit vectors $\mathbf{x}$ with $s \rightarrow 0$ for which $|\mathbf{x}^*\mathbf{A}\mathbf{x} - \lambda_1|$ is **not** $O(s^2)$.
- Hence say which hypothesis of [Theorem 24.4.4](#) the example violates, and which step of the proof breaks.

Solution.

- This is Step 1 of the proof of [Theorem 24.4.4](#) with $j = 1$. Expanding $\mathbf{x} = \sum_i c_i \mathbf{q}_i$ with $\sum_i |c_i|^2 = 1$, we get $s^2 = \sum_{i \neq 1} |c_i|^2$ and, by [Theorem 11.2.1 \(b\)](#), $R_{\mathbf{A}}(\mathbf{x}) = \sum_i |c_i|^2 \lambda_i$.

Since the weights sum to 1,

$$R_{\mathbf{A}}(\mathbf{x}) - \lambda_1 = \sum_{i \neq 1} |c_i|^2 (\lambda_i - \lambda_1),$$

whose modulus is at most $s^2 \max_{i \neq 1} |\lambda_i - \lambda_1|$.

- (b) Take $\mathbf{A} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$, the standard nondiagonalizable example, with $\lambda_1 = 0$ and $\mathbf{q}_1 = \mathbf{e}_1$. Let $\mathbf{x} = (\sqrt{1-s^2}, s)$ for small $s > 0$, a unit vector whose part orthogonal to $\mathbf{e}_1$ has length exactly s . Then

$$\mathbf{x}^* \mathbf{A} \mathbf{x} = \overline{x_1} x_2 = s \sqrt{1-s^2},$$

which is $s + O(s^3)$: **first** order in s , not second. So no bound of the form Cs^2 can hold as $s \rightarrow 0$.

- (c) The example violates the hypothesis that $\mathbf{A}$ be Hermitian. The step that breaks is Step 1: for a Hermitian $\mathbf{A}$ the quantity $\mathbf{x}^* \mathbf{A} \mathbf{x}$ is the weighted average $\sum_i |c_i|^2 \lambda_i$ of the eigenvalues, in which only the squared moduli of the coefficients appear and the first-order term is therefore absent. For a general $\mathbf{A}$ the cross terms $\overline{c_i} c_k \mathbf{q}_i^* \mathbf{A} \mathbf{q}_k$ with $i \neq k$ survive, and they are of first order in the small coefficients. With Step 1 weakened to $O(s)$, Step 3 multiplies s by $O(s)$ and the rate is quadratic, not cubic — which is the remark after [Theorem 24.4.4](#).

Exercise. C2: Two eigenvalues of equal modulus, over the reals

Let $\mathbf{A} \in M_2(\mathbb{R})$ be the rotation matrix $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$.

- Find the eigenvalues of $\mathbf{A}$ and explain why [Theorem 24.4.1](#) does not apply.
- Describe the power iteration from an arbitrary unit $\mathbf{x}^{(0)} \in \mathbb{R}^2$, and show that no line is a limit.
- Now take the shift $\mu = \frac{1}{2}$ and run inverse iteration over $\mathbb{C}$. Does the theorem apply? Compute ρ_μ . *Hint: the eigenvalues of $\mathbf{A}$ are $\pm i$.*
- Deduce a one-sentence moral about which of the two failures of [Example \(The method that never converges\)](#) a shift can repair.

Solution.

- $p_{\mathbf{A}}(x) = x^2 + 1$, so the eigenvalues are i and $-i$, both of modulus 1. The hypothesis $|\lambda_1| > |\lambda_2|$ fails.
- $\mathbf{A}$ is orthogonal, so $\|\mathbf{A}\mathbf{x}\| = \|\mathbf{x}\|$ and no normalization is needed: $\mathbf{x}^{(k)} = \mathbf{A}^k \mathbf{x}^{(0)}$. Since $\mathbf{A}^2 = -\mathbf{I}$, the iterates are $\mathbf{x}^{(0)}, \mathbf{A}\mathbf{x}^{(0)}, -\mathbf{x}^{(0)}, -\mathbf{A}\mathbf{x}^{(0)}, \mathbf{x}^{(0)}, \dots$, of period 4, and the lines they span alternate between $\text{span}(\mathbf{x}^{(0)})$ and $\text{span}(\mathbf{A}\mathbf{x}^{(0)})$, which are distinct because $\mathbf{A}$ has no real eigenvector. A sequence taking two distinct values infinitely often has no limit.
- Over $\mathbb{C}$ the matrix is diagonalizable, with eigenvectors $(1, i)$ and $(1, -i)$. The shift $\mu = \frac{1}{2}$ is not an eigenvalue, and the distances are $|i - \frac{1}{2}| = |-i - \frac{1}{2}| = \sqrt{5}/2$ — **equal**. So there is no unique nearest eigenvalue and [Theorem 24.4.3](#) still does not apply; $\rho_\mu = 1$. Any real shift fails in the same way, by symmetry in the conjugate pair; a shift such as $\mu = \frac{i}{2}$ does work, with distances $\frac{1}{2}$ and $\frac{3}{2}$ and $\rho_\mu = \frac{1}{3}$.
- A shift repairs a bad *rate* and a wrong *target*, by moving the origin from which moduli are measured; it cannot repair a tie, and over $\mathbb{R}$ a conjugate pair of eigenvalues is a tie

that no real shift can break.

Exercise. C3: A small residual certifies an eigenpair

Let $\mathbf{A} \in M_n(\mathbb{C})$ be Hermitian with orthonormal eigenbasis $\mathbf{q}_1, \dots, \mathbf{q}_n$, let $\hat{\mathbf{q}}$ be a unit vector, let $\hat{\mu} = R_{\mathbf{A}}(\hat{\mathbf{q}})$ and $\varepsilon = \|\mathbf{A}\hat{\mathbf{q}} - \hat{\mu}\hat{\mathbf{q}}\|$. Suppose λ_j is the only eigenvalue of $\mathbf{A}$ in the interval $(\hat{\mu} - g, \hat{\mu} + g)$.

- Prove that $\|\hat{\mathbf{q}} - \mathbf{q}_j \mathbf{q}_j^* \hat{\mathbf{q}}\| \leq \varepsilon/g$, assuming $\varepsilon < g$.
- Deduce a bound on $|\hat{\mu} - \lambda_j|$ from Theorem 20.6.2, and compare it with the bound from Theorem 20.6.1.
- In the inverse iteration of Definition 24.4.2 with a backward stable solve, ε is about $1/\|\hat{\mathbf{z}}\| + cu\|\mathbf{A} - \mu\mathbf{I}\|$. Explain what (a) then says about the computed direction, and in particular why the accuracy improves as the shift improves.

Solution.

- Expand $\hat{\mathbf{q}} = \sum_i d_i \mathbf{q}_i$ with $\sum_i |d_i|^2 = 1$. Then $\mathbf{A}\hat{\mathbf{q}} - \hat{\mu}\hat{\mathbf{q}} = \sum_i d_i(\lambda_i - \hat{\mu})\mathbf{q}_i$, so by Theorem 11.2.1 (c),

$$\varepsilon^2 = \sum_i |d_i|^2 (\lambda_i - \hat{\mu})^2 \geq \sum_{i \neq j} |d_i|^2 (\lambda_i - \hat{\mu})^2 \geq g^2 \sum_{i \neq j} |d_i|^2,$$

because every λ_i with $i \neq j$ lies outside $(\hat{\mu} - g, \hat{\mu} + g)$. Since $\|\hat{\mathbf{q}} - \mathbf{q}_j \mathbf{q}_j^* \hat{\mathbf{q}}\|^2 = \sum_{i \neq j} |d_i|^2$, taking square roots gives the bound.

- Theorem 20.6.1 gives an eigenvalue within ε of $\hat{\mu}$, and by hypothesis the only candidate in that range (once $\varepsilon < g$) is λ_j , so $|\hat{\mu} - \lambda_j| \leq \varepsilon$. Theorem 20.6.2, applied with $\alpha = \hat{\mu} - g$ and $\beta = \hat{\mu} + g$, improves this to

$$|\hat{\mu} - \lambda_j| \leq \frac{\varepsilon^2}{g},$$

which is smaller by a factor ε/g — the quadratic gain of the Rayleigh quotient once more.

- It says that the **direction** of the computed $\hat{\mathbf{z}}$ is within ε/g of $\mathbf{q}_j$, whatever the relative error in $\hat{\mathbf{z}}$ itself. As the shift μ approaches λ_j , the matrix $\mathbf{A} - \mu\mathbf{I}$ becomes more nearly singular, and — provided the right-hand side $\mathbf{x}^{(k)}$ has a component along $\mathbf{q}_j$ that is not small, by the bound $\|\hat{\mathbf{z}}\| \geq |\mathbf{q}_j^* \mathbf{x}^{(k)}| / (|\lambda_j - \mu| + \|\Delta\|)$ established in the discussion following Theorem 24.4.3 — $\|\hat{\mathbf{z}}\|$ grows like $1/|\lambda_j - \mu|$. So the first part of ε shrinks in proportion, and the second part is already at the rounding level. **A worse-conditioned system gives a better answer**, because the quantity that grows is the one in the denominator. What the bound never forgives is a small g : two nearby eigenvalues leave ε/g large no matter how good the shift is.