

Math 6610: Analysis of Numerical Methods, I

Rayleigh iteration and the QR algorithm with shifts

Department of Mathematics, University of Utah

Fall 2025

Resources: Trefethen and Bau 1997, Lectures 27, 29
Atkinson 1989, Section 9.6
Süli and Mayers 2003, Sections 5.8, 5.9
Salgado and Wise 2022, Section 8.4

Recall: power iteration

With $\mathbf{A} \in \mathbb{C}^{n \times n}$ and $\mathbf{x} \in \mathbb{C}^n$, we expect

$$\mathbf{A}^k \mathbf{x} \propto \lambda_1^k \mathbf{v}_1, \quad k \gg 1,$$

where $(\lambda_1, \mathbf{v}_1)$ is the dominant eigenpair of \mathbf{A} . The generic algorithm that computes $\mathbf{A}^k \mathbf{x}$ as an approximation to \mathbf{v}_1 is power iteration.

Convergence is *linear*, i.e., $\|\mathbf{A}^k \mathbf{x} - c\lambda_1^k \mathbf{v}_1\|_2 \sim r^k$, where $r = \sup_{j>1} |\lambda_j/\lambda_1| < 1$.

To compute several eigenpairs, we can use simultaneous (power) iteration, computing \mathbf{A}^k for large k .

A *stable* numerical algorithm that essentially performs power iteration is the QR algorithm.

Recall: power iteration

With $\mathbf{A} \in \mathbb{C}^{n \times n}$ and $\mathbf{x} \in \mathbb{C}^n$, we expect

$$\mathbf{A}^k \mathbf{x} \propto \lambda_1^k \mathbf{v}_1, \quad k \gg 1,$$

where $(\lambda_1, \mathbf{v}_1)$ is the dominant eigenpair of \mathbf{A} . The generic algorithm that computes $\mathbf{A}^k \mathbf{x}$ as an approximation to \mathbf{v}_1 is power iteration.

Convergence is *linear*, i.e., $\|\mathbf{A}^k \mathbf{x} - c\lambda_1^k \mathbf{v}_1\|_2 \sim r^k$, where $r = \sup_{j>1} |\lambda_j/\lambda_1| < 1$.

To compute several eigenpairs, we can use simultaneous (power) iteration, computing \mathbf{A}^k for large k .

A *stable* numerical algorithm that essentially performs power iteration is the QR algorithm.

The problem: The constant r above can be very close to 1, making convergence slow, as then $r^k \approx 1 - k(1 - r)$.

Here's a simple but powerful observation: if λ_j is an eigenvalue of \mathbf{A} , and $\mu \in \mathbb{C} \setminus \lambda(\mathbf{A})$, then

$$\frac{1}{\lambda_j - \mu} \in \lambda(\mathbf{B}), \quad \mathbf{B} = (\mathbf{A} - \mu \mathbf{I})^{-1},$$

Here's a simple but powerful observation: if λ_j is an eigenvalue of \mathbf{A} , and $\mu \in \mathbb{C} \setminus \lambda(\mathbf{A})$, then

$$\frac{1}{\lambda_j - \mu} \in \lambda(\mathbf{B}), \quad \mathbf{B} = (\mathbf{A} - \mu \mathbf{I})^{-1},$$

The main utility of this fact is that if $\mu \approx \lambda_j$, then

$$\max_{k \neq j} \frac{\frac{1}{\lambda_k - \mu}}{\frac{1}{\lambda_j - \mu}} = \max_{k \neq j} \frac{\lambda_j - \mu}{\lambda_k - \mu} \ll 1.$$

The factor on the left is constant “ r ” in the convergence rate of power iteration for $(\mathbf{A} - \mu \mathbf{I})^{-1}$.

Here's a simple but powerful observation: if λ_j is an eigenvalue of \mathbf{A} , and $\mu \in \mathbb{C} \setminus \lambda(\mathbf{A})$, then

$$\frac{1}{\lambda_j - \mu} \in \lambda(\mathbf{B}), \quad \mathbf{B} = (\mathbf{A} - \mu \mathbf{I})^{-1},$$

The main utility of this fact is that if $\mu \approx \lambda_j$, then

$$\max_{k \neq j} \frac{\frac{1}{\lambda_k - \mu}}{\frac{1}{\lambda_j - \mu}} = \max_{k \neq j} \frac{\lambda_j - \mu}{\lambda_k - \mu} \ll 1.$$

The factor on the left is constant “ r ” in the convergence rate of power iteration for $(\mathbf{A} - \mu \mathbf{I})^{-1}$.

The punch line: if we can choose $\mu \approx \lambda_j$, we can (considerably) accelerate power iteration's convergence to λ_j .

The parameter μ is called a *shift*.

To see how shifting can provide utility, note that even if $|\lambda_1 - \lambda_2|$ is small, if we are able to choose μ very close to λ_1 , then

$$\left| \frac{\lambda_1 - \mu}{\lambda_2 - \mu} \right| \ll 1, \quad \left| \frac{1}{\lambda_1 - \mu} \right| \gg \left| \frac{1}{\lambda_2 - \mu} \right|,$$

and hence power iteration on $(\mathbf{A} - \mu\mathbf{I})$ has a well-separated dominant eigenvalue of $1/(\lambda_1 - \mu)$, which can be inverted and re-shifted back to λ_1 since μ is known.

To see how shifting can provide utility, note that even if $|\lambda_1 - \lambda_2|$ is small, if we are able to choose μ very close to λ_1 , then

$$\left| \frac{\lambda_1 - \mu}{\lambda_2 - \mu} \right| \ll 1, \quad \left| \frac{1}{\lambda_1 - \mu} \right| \gg \left| \frac{1}{\lambda_2 - \mu} \right|,$$

and hence power iteration on $(\mathbf{A} - \mu\mathbf{I})$ has a well-separated dominant eigenvalue of $1/(\lambda_1 - \mu)$, which can be inverted and re-shifted back to λ_1 since μ is known.

Performing power iteration on $(\mathbf{A} - \mu\mathbf{I})^{-1}$ is called (shifted) inverse iteration.

0. Given $\mu \in \mathbb{C}$, define $\mathbf{A}^{\mu^-} = (\mathbf{A} - \mu\mathbf{I})^{-1}$. Randomly initialize \mathbf{x}_0 , set $j = 0$.
1. Compute $\mathbf{x}_{j+1} = \frac{\mathbf{A}^{\mu^-} \mathbf{x}_j}{\|\mathbf{A}^{\mu^-} \mathbf{x}_j\|_2}$.
2. Compute $\mu_{j+1} = R_{\mathbf{A}}(\mathbf{x}_{j+1})$.
3. Set $j \leftarrow j + 1$, return to step 1.

NB: This requires linear solves!

Note that the choice of shift is integral to the success of inverse iteration, *and* the choice of shift should depend on the “targeted” eigenvalue.

Happily, we have a reasonable guess for a good eigenvalue: it's the Rayleigh quotient $R_A(\mathbf{x}_j)$.

Note that the choice of shift is integral to the success of inverse iteration, *and* the choice of shift should depend on the “targeted” eigenvalue.

Happily, we have a reasonable guess for a good eigenvalue: it's the Rayleigh quotient $R_A(\mathbf{x}_j)$.

Therefore, one strategy is to actually use this estimate of the eigenvalue as a shift:

$$\mathbf{x}_0 \rightarrow \mu = R_A(\mathbf{x}_0) \rightarrow \mathbf{x}_1 = \frac{\mathbf{A}^{\mu_0} \mathbf{x}_0}{\|\mathbf{A}^{\mu_0} \mathbf{x}_0\|} \rightarrow \mu_1 = R_A(\mathbf{x}_1) \rightarrow \dots$$

This leads to an algorithm called Rayleigh iteration, which uses the Rayleigh quotient to approximate shifts for use in inverse iteration:

0. With (μ_0, \mathbf{x}_0) a guess for an initial eigenpair, set $j = 0$.
1. Compute $\mathbf{x}_{j+1} = \frac{\mathbf{A}^{\mu_j} \mathbf{x}_j}{\|\mathbf{A}^{\mu_j} \mathbf{x}_j\|_2}$.
2. Compute $\mu_{j+1} = R_{\mathbf{A}}(\mathbf{x}_{j+1})$.
3. Set $j \leftarrow j + 1$, return to step 1.

(Alternatively, just initialize \mathbf{x}_0 and set $\mu_0 = R_{\mathbf{A}}(\mathbf{x}_0)$).

This leads to an algorithm called Rayleigh iteration, which uses the Rayleigh quotient to approximate shifts for use in inverse iteration:

0. With (μ_0, \mathbf{x}_0) a guess for an initial eigenpair, set $j = 0$.
1. Compute $\mathbf{x}_{j+1} = \frac{\mathbf{A}^{\mu_j} \mathbf{x}_j}{\|\mathbf{A}^{\mu_j} \mathbf{x}_j\|_2}$.
2. Compute $\mu_{j+1} = R_{\mathbf{A}}(\mathbf{x}_{j+1})$.
3. Set $j \leftarrow j + 1$, return to step 1.

(Alternatively, just initialize \mathbf{x}_0 and set $\mu_0 = R_{\mathbf{A}}(\mathbf{x}_0)$. How well does Rayleigh iteration work? Consider first power iteration with (λ, \mathbf{v}) a dominant eigenpair:

$$\mathbf{x}_0 = \mathbf{v} + \boldsymbol{\delta}, \quad \left\| \frac{\mathbf{A}}{\lambda} \boldsymbol{\delta} \right\|_2 \leq r \|\boldsymbol{\delta}\|_2, \quad r = \max_{\lambda \in \Lambda(\mathbf{A}) \setminus \{\lambda\}} \left| \frac{\lambda_j}{\lambda} \right| < 1,$$

This leads to an algorithm called Rayleigh iteration, which uses the Rayleigh quotient to approximate shifts for use in inverse iteration:

0. With (μ_0, \mathbf{x}_0) a guess for an initial eigenpair, set $j = 0$.
1. Compute $\mathbf{x}_{j+1} = \frac{\mathbf{A}^{\mu_j} \mathbf{x}_j}{\|\mathbf{A}^{\mu_j} \mathbf{x}_j\|_2}$.
2. Compute $\mu_{j+1} = R_{\mathbf{A}}(\mathbf{x}_{j+1})$.
3. Set $j \leftarrow j + 1$, return to step 1.

(Alternatively, just initialize \mathbf{x}_0 and set $\mu_0 = R_{\mathbf{A}}(\mathbf{x}_0)$. How well does Rayleigh iteration work? Consider first power iteration with (λ, \mathbf{v}) a dominant eigenpair:

$$\mathbf{x}_0 = \mathbf{v} + \boldsymbol{\delta}, \quad \left\| \frac{\mathbf{A}}{\lambda} \boldsymbol{\delta} \right\|_2 \leq r \|\boldsymbol{\delta}\|_2, \quad r = \max_{\lambda \in \Lambda(\mathbf{A}) \setminus \{\lambda\}} \left| \frac{\lambda_j}{\lambda} \right| < 1,$$

where we've assumed \mathbf{A} is diagonalizable. Then,

$$\left\| \frac{1}{\lambda^k} \mathbf{A}^k \mathbf{x}_0 - \mathbf{v} \right\|_2 = \left\| \frac{\mathbf{A}}{\lambda} \underbrace{\left(\frac{\mathbf{A}^{k-1} \mathbf{x}_0}{\lambda^{k-1}} - \mathbf{v} \right)}_{\propto \boldsymbol{\delta}} \right\|_2 \leq r \left\| \frac{1}{\lambda^{k-1}} \mathbf{A}^{k-1} \mathbf{x}_0 - \mathbf{v} \right\|_2.$$

I.e., the error at stage k is essentially the error at stage $k - 1$ times the factor $r < 1$.

$$\left\| \frac{1}{\lambda_1^k} \mathbf{A}^k \mathbf{x}_0 - \mathbf{v} \right\|_2 \leq r \left\| \frac{1}{\lambda_1^{k-1}} \mathbf{A}^{k-1} \mathbf{x}_0 - \mathbf{v} \right\|_2 \implies \|\mathbf{x}_{k+1} - \mathbf{v}\|_2 \leq r \|\mathbf{x}_k - \mathbf{v}\|_2$$

The next part we need to analyze is the Rayleigh quotient, which is how we compute shifts for the next iteration.

Recall that if (λ, \mathbf{v}) is an eigenpair of \mathbf{A} , then $R_{\mathbf{A}}(\mathbf{v}) = \lambda$.

$$\left\| \frac{1}{\lambda_1^k} \mathbf{A}^k \mathbf{x}_0 - \mathbf{v} \right\|_2 \leq r \left\| \frac{1}{\lambda_1^{k-1}} \mathbf{A}^{k-1} \mathbf{x}_0 - \mathbf{v} \right\|_2 \implies \|\mathbf{x}_{k+1} - \mathbf{v}\|_2 \leq r \|\mathbf{x}_k - \mathbf{v}\|$$

The next part we need to analyze is the Rayleigh quotient, which is how we compute shifts for the next iteration.

Recall that if (λ, \mathbf{v}) is an eigenpair of \mathbf{A} , then $R_{\mathbf{A}}(\mathbf{v}) = \lambda$.

A nice additional property is that $(\nabla_{\mathbf{x}} R_{\mathbf{A}}(\mathbf{x}))|_{\mathbf{x}=\mathbf{v}} = 0$.
I.e., eigenvectors \mathbf{v} are stationary points of $R_{\mathbf{A}}$.

Hence, by a Taylor series argument, when $\mathbf{x} \approx \mathbf{v}$, we have,

$$R_{\mathbf{A}}(\mathbf{x}) = R_{\mathbf{A}}(\mathbf{v}) + \mathcal{O}(\|\mathbf{x} - \mathbf{v}\|_2^2) = \lambda + \mathcal{O}(\|\mathbf{x} - \mathbf{v}\|_2^2).$$

$$\left\| \frac{1}{\lambda_1^k} \mathbf{A}^k \mathbf{x}_0 - \mathbf{v} \right\|_2 \leq r \left\| \frac{1}{\lambda^{k-1}} \mathbf{A}^{k-1} \mathbf{x}_0 - \mathbf{v} \right\|_2 \implies \|\mathbf{x}_{k+1} - \mathbf{v}\|_2 \leq r \|\mathbf{x}_k - \mathbf{v}\|$$

The next part we need to analyze is the Rayleigh quotient, which is how we compute shifts for the next iteration.

Recall that if (λ, \mathbf{v}) is an eigenpair of \mathbf{A} , then $R_{\mathbf{A}}(\mathbf{v}) = \lambda$.

A nice additional property is that $(\nabla_{\mathbf{x}} R_{\mathbf{A}}(\mathbf{x}))|_{\mathbf{x}=\mathbf{v}} = 0$.
I.e., eigenvectors \mathbf{v} are stationary points of $R_{\mathbf{A}}$.

Hence, by a Taylor series argument, when $\mathbf{x} \approx \mathbf{v}$, we have,

$$R_{\mathbf{A}}(\mathbf{x}) = R_{\mathbf{A}}(\mathbf{v}) + \mathcal{O}(\|\mathbf{x} - \mathbf{v}\|_2^2) = \lambda + \mathcal{O}(\|\mathbf{x} - \mathbf{v}\|_2^2).$$

Therefore, if (μ_k, \mathbf{x}_k) is “close” to an eigenpair (λ, \mathbf{v}) of \mathbf{A} , then $1/(\lambda - \mu)$ is the dominant eigenpair of \mathbf{A}_{μ^-} , so,

$$\begin{aligned} r &= \max_{\lambda \in \Lambda(\mathbf{A}) \setminus \{\lambda\}} \frac{\mu_k - \lambda}{\mu_k - \lambda} \\ &= \mathcal{O}(\mu_k - \lambda) = \mathcal{O}(R_{\mathbf{A}}(\mathbf{x}_k) - R_{\mathbf{A}}(\mathbf{v})) = \mathcal{O}(\|\mathbf{x}_k - \mathbf{v}\|_2^2). \end{aligned}$$

We have “approximately” proved the following:

Theorem

Let $\mathbf{A} \in \mathbb{C}^{n \times n}$. For almost every initialization $(\mu_0, \mathbf{x}_0) \in \mathbb{C} \times \mathbb{C}^n$ of Rayleigh iteration, there is some eigenpair (λ, \mathbf{v}) of \mathbf{A} such that,

$$\|\mathbf{x}_{j+1} - \mathbf{v}\|_2 = \mathcal{O}(\|\mathbf{x}_j - \mathbf{v}\|_2^3),$$

$$\|\mu_{j+1} - \lambda\|_2 = \mathcal{O}(\|\mu_j - \lambda\|_2^3)$$

We have “approximately” proved the following:

Theorem

Let $\mathbf{A} \in \mathbb{C}^{n \times n}$. For almost every initialization $(\mu_0, \mathbf{x}_0) \in \mathbb{C} \times \mathbb{C}^n$ of Rayleigh iteration, there is some eigenpair (λ, \mathbf{v}) of \mathbf{A} such that,

$$\|\mathbf{x}_{j+1} - \mathbf{v}\|_2 = \mathcal{O}(\|\mathbf{x}_j - \mathbf{v}\|_2^3), \quad \|\mu_{j+1} - \lambda\|_2 = \mathcal{O}(\|\mu_j - \lambda\|_2^3)$$

The result: convergence is cubic. (!!!!)

I.e., if we reach a stage where (μ_j, \mathbf{v}_j) has one digit of accuracy, then $(\mu_{j+1}, \mathbf{v}_{j+1})$ will have 3 digits of accuracy, and $(\mu_{j+2}, \mathbf{v}_{j+2})$ will have *nine*. (!!!!)

We have “approximately” proved the following:

Theorem

Let $\mathbf{A} \in \mathbb{C}^{n \times n}$. For almost every initialization $(\mu_0, \mathbf{x}_0) \in \mathbb{C} \times \mathbb{C}^n$ of Rayleigh iteration, there is some eigenpair (λ, \mathbf{v}) of \mathbf{A} such that,

$$\|\mathbf{x}_{j+1} - \mathbf{v}\|_2 = \mathcal{O}(\|\mathbf{x}_j - \mathbf{v}\|_2^3), \quad \|\mu_{j+1} - \lambda\|_2 = \mathcal{O}(\|\mu_j - \lambda\|_2^3)$$

The result: convergence is cubic. (!!!!)

I.e., if we reach a stage where (μ_j, \mathbf{v}_j) has one digit of accuracy, then $(\mu_{j+1}, \mathbf{v}_{j+1})$ will have 3 digits of accuracy, and $(\mu_{j+2}, \mathbf{v}_{j+2})$ will have *nine*. (!!!!)

Hence, Rayleigh iteration is an *exceptional* algorithm for computing the spectrum.

Paired with deflation, one can compute the whole spectrum very efficiently.

A “Rayleigh” version of the QR algorithm?

Power iteration was our first algorithm for computing eigenvalues.

The QR algorithm was a stable way to perform (simultaneous) power iteration.

We now have a “better” algorithm, Rayleigh iteration.

Is there a stable way to perform inverse iteration (+shifts), like the QR algorithm?

Recall the (“pure”) QR algorithm with $A_0 = A$:

$$A_j \stackrel{\text{QR}}{=} Q_j R_j,$$

$$A_{j+1} = R_j Q_j = Q_j^* A_j Q_j,$$

and that,

$$A_k = (Q^{(k)})^* A Q^{(k)},$$

$$A^k \stackrel{\text{QR}}{=} Q^{(k)} R^{(k)}$$

Recall the (“pure”) QR algorithm with $A_0 = A$:

$$A_j \stackrel{\text{QR}}{=} Q_j R_j, \quad A_{j+1} = R_j Q_j = Q_j^* A_j Q_j,$$

and that,

$$A_k = (Q^{(k)})^* A Q^{(k)}, \quad A^k \stackrel{\text{QR}}{=} Q^{(k)} R^{(k)}$$

By taking inverses, we can actually see how this is related to inverse iteration:

$$A^{-k} = (R^{(k)})^{-1} (Q^{(k)})^* \implies (A^{-T})^k = Q_*^{(k)} (R^{(k)})^{-T}$$

where $Q_*^{(k)}$ is the componentwise complex conjugate of $Q^{(k)}$, and like $Q^{(k)}$ is unitary.

We now introduce a permutation matrix $\mathbf{\Pi} \in \mathbb{C}^{n \times n}$ with ones on the antidiagonal: $(\mathbf{\Pi})_{j, n-j+1} = 1$.

I.e., $\mathbf{\Pi}\mathbf{x}$ reverses (geometrically flips) the entries of \mathbf{x} .

$$\mathbf{\Pi}\mathbf{x} = (x_n, x_{n-1}, \dots, x_2, x_1).$$

And we also have that $\mathbf{\Pi}^{-1} = \mathbf{\Pi}^* = \mathbf{\Pi}$, i.e., $\mathbf{\Pi}^2 = \mathbf{I}$.

We now introduce a permutation matrix $\mathbf{\Pi} \in \mathbb{C}^{n \times n}$ with ones on the antidiagonal: $(\mathbf{\Pi})_{j, n-j+1} = 1$.

I.e., $\mathbf{\Pi}\mathbf{x}$ reverses (geometrically flips) the entries of \mathbf{x} .

$$\mathbf{\Pi}\mathbf{x} = (x_n, x_{n-1}, \dots, x_2, x_1).$$

And we also have that $\mathbf{\Pi}^{-1} = \mathbf{\Pi}^* = \mathbf{\Pi}$, i.e., $\mathbf{\Pi}^2 = \mathbf{I}$.

With this matrix, we have:

$$(\mathbf{A}^{-T})^k \mathbf{\Pi} = \underbrace{\mathbf{Q}_*^{(k)}}_{\text{unitary}} \underbrace{\mathbf{\Pi} \mathbf{\Pi} (\mathbf{R}^{(k)})^{-T} \mathbf{\Pi}}_{\text{upper triangular}}$$

We now introduce a permutation matrix $\mathbf{\Pi} \in \mathbb{C}^{n \times n}$ with ones on the antidiagonal: $(\mathbf{\Pi})_{j, n-j+1} = 1$.

i.e., $\mathbf{\Pi}\mathbf{x}$ reverses (geometrically flips) the entries of \mathbf{x} .

$$\mathbf{\Pi}\mathbf{x} = (x_n, x_{n-1}, \dots, x_2, x_1).$$

And we also have that $\mathbf{\Pi}^{-1} = \mathbf{\Pi}^* = \mathbf{\Pi}$, i.e., $\mathbf{\Pi}^2 = \mathbf{I}$.

With this matrix, we have:

$$(\mathbf{A}^{-T})^k \mathbf{\Pi} = \underbrace{\mathbf{Q}_*^{(k)}}_{\text{unitary}} \underbrace{\mathbf{\Pi} \mathbf{\Pi} (\mathbf{R}^{(k)})^{-T} \mathbf{\Pi}}_{\text{upper triangular}}$$

i.e., the pure QR algorithm actually computes the Q factor for \mathbf{A}^{-T} , i.e., it performs inverse iteration as well!

(More precisely, the componentwise conjugate of $\mathbf{Q}^{(k)}$, with its columns reversed, is the same as the Q factor for inverse iteration on \mathbf{A}^T .)

In particular, the last column of $\mathbf{Q}^{(k)}$ has the (conjugated) first column of inverse iteration on \mathbf{A}^T .

We have learned that the QR algorithm is both power iteration and (essentially) also inverse iteration.

How do we build in shifts? Suppose that μ_1, μ_2, \dots , are *a priori* known shifts. Consider the following *shifted* QR algorithm:

0. Let $j = 1$, $\mathbf{A}_1 = \mathbf{A}$.
1. Define $\mathbf{A}_j - \mu_j \mathbf{I} \stackrel{\text{QR}}{=} \mathbf{Q}_j \mathbf{R}_j$.
2. Compute $\mathbf{A}_{j+1} := \mathbf{R}_j \mathbf{Q}_j + \mu_j \mathbf{I}$
3. Set $j \leftarrow j + 1$, return to step 1.

We have learned that the QR algorithm is both power iteration and (essentially) also inverse iteration.

How do we build in shifts? Suppose that μ_1, μ_2, \dots , are *a priori* known shifts. Consider the following *shifted* QR algorithm:

0. Let $j = 1$, $\mathbf{A}_1 = \mathbf{A}$.
1. Define $\mathbf{A}_j - \mu_j \mathbf{I} \stackrel{\text{QR}}{=} \mathbf{Q}_j \mathbf{R}_j$.
2. Compute $\mathbf{A}_{j+1} := \mathbf{R}_j \mathbf{Q}_j + \mu_j \mathbf{I}$
3. Set $j \leftarrow j + 1$, return to step 1.

Some exercises, similar to what is done for the pure QR algorithm, reveal:

$$\mathbf{A}_{j+1} = \mathbf{Q}_j^* \mathbf{A}_j \mathbf{Q}_j, \quad \mathbf{Q}^{(k)} = \prod_{j=1}^k \mathbf{Q}_j = \mathbf{Q}_1 \cdots \mathbf{Q}_k, \quad \mathbf{R}^{(k)} = \prod_{j=k}^1 \mathbf{R}_j = \mathbf{R}_k \cdots \mathbf{R}_1,$$

where $\mathbf{Q}^{(k)}$ and $\mathbf{R}^{(k)}$ satisfy,

$$\prod_{j=k}^1 (\mathbf{A} - \mu_j \mathbf{I}) = (\mathbf{A} - \mu_k \mathbf{I}) \cdots (\mathbf{A} - \mu_1 \mathbf{I}) = \mathbf{Q}^{(k)} \mathbf{R}^{(k)},$$

We have established that:

- The QR algorithm perform inverse iteration
- The shifted QR algorithm therefore performs shifted inverse iteration (on a transpose of A)

Recall that our goal, Rayleigh iteration, simply defines the shifts through the Rayleigh quotient. How do we compute Rayleigh quotients?

We have established that:

- The QR algorithm perform inverse iteration
- The shifted QR algorithm therefore performs shifted inverse iteration (on a transpose of \mathbf{A})

Recall that our goal, Rayleigh iteration, simply defines the shifts through the Rayleigh quotient. How do we compute Rayleigh quotients?

The somewhat surprising answer: we essentially get them for free:

$$(A_k)_{n,n} = \mathbf{e}_n^* \mathbf{A}_k \mathbf{e}_n = \mathbf{e}_n^* (\mathbf{Q}^{(k)})^* \mathbf{A} \mathbf{Q}^{(k)} \mathbf{e}_n = \mathbf{e}_n^T (\mathbf{Q}^{(k)})^T \mathbf{A}^T (\mathbf{Q}^{(k)})_* \mathbf{e}_n = R_{\mathbf{A}^T}(\mathbf{y}),$$

where \mathbf{y} is the last column of the componentwise conjugate of $\mathbf{Q}^{(k)}$, which is precisely shifted inverse iteration's estimate of the dominant eigenvector.

The QR algorithm with shifts

The following is the QR algorithm with *Rayleigh quotient* shifts:

0. Let $j = 1$, $\mathbf{A}_1 = \mathbf{A}$.
1. Define $\mu_j = (A_j)_{n,n}$.
2. Define $\mathbf{A}_j - \mu_j \mathbf{I} \stackrel{\text{QR}}{=} \mathbf{Q}_j \mathbf{R}_j$.
3. Compute $\mathbf{A}_{j+1} := \mathbf{R}_j \mathbf{Q}_j + \mu_j \mathbf{I}$
4. Set $j \leftarrow j + 1$, return to step 1.

This algorithm performs (very well-disguised) Rayleigh iteration.

The following is the QR algorithm with *Rayleigh quotient* shifts:

0. Let $j = 1$, $\mathbf{A}_1 = \mathbf{A}$.
1. Define $\mu_j = (A_j)_{n,n}$.
2. Define $\mathbf{A}_j - \mu_j \mathbf{I} \stackrel{\text{QR}}{=} \mathbf{Q}_j \mathbf{R}_j$.
3. Compute $\mathbf{A}_{j+1} := \mathbf{R}_j \mathbf{Q}_j + \mu_j \mathbf{I}$
4. Set $j \leftarrow j + 1$, return to step 1.

This algorithm performs (very well-disguised) Rayleigh iteration.

This algorithm can fail (to converge) if the shift happens to be equidistant between dominant eigenvalues. There are other (slightly more complicated) choices of shifts that are much more robust.

The following is the QR algorithm with *Rayleigh quotient* shifts:

0. Let $j = 1$, $\mathbf{A}_1 = \mathbf{A}$.
1. Define $\mu_j = (A_j)_{n,n}$.
2. Define $\mathbf{A}_j - \mu_j \mathbf{I} \stackrel{\text{QR}}{=} \mathbf{Q}_j \mathbf{R}_j$.
3. Compute $\mathbf{A}_{j+1} := \mathbf{R}_j \mathbf{Q}_j + \mu_j \mathbf{I}$
4. Set $j \leftarrow j + 1$, return to step 1.

This algorithm performs (very well-disguised) Rayleigh iteration.

This algorithm can fail (to converge) if the shift happens to be equidistant between dominant eigenvalues. There are other (slightly more complicated) choices of shifts that are much more robust.

This algorithm requires deflation for practicality (otherwise everything continues to focus on the dominant eigenvalue).

When the dominant eigenvalue is found, then the last row of \mathbf{A}_j is zero, except for the last entry (which contains the eigenvalue).





To deflate, one may slice the last row + column off and continue.

The QR algorithm (with wisely chosen shifts) was the gold standard for eigenvalue computation in the 20th century.

Recently, new types of algorithms (divide-and-conquer algorithms and relatively robust representation, RRR, algorithms) are also used.

The main advantage of these alternatives was typically in either complexity or storage, but the gains are not in terms of asymptotic behavior, and instead in terms of decreasing the constants in those asymptotic behaviors.

Nevertheless, LAPACK, the set of routines used nearly ubiquitously for numerical computing, still has polished QR-with-shifts eigenvalue algorithms that are still used.

-  Atkinson, Kendall (1989). *An Introduction to Numerical Analysis*. New York: Wiley. ISBN: 978-0-471-62489-9.
-  Salgado, Abner J. and Steven M. Wise (2022). *Classical Numerical Analysis: A Comprehensive Course*. Cambridge: Cambridge University Press. ISBN: 978-1-108-83770-5. DOI: 10.1017/9781108942607.
-  Süli, Endre and David F. Mayers (2003). *An Introduction to Numerical Analysis*. Cambridge: Cambridge University Press. ISBN: 978-0-521-00794-8. DOI: 10.1017/CB09780511801181.
-  Trefethen, Lloyd N. and David Bau (1997). *Numerical Linear Algebra*. SIAM: Society for Industrial and Applied Mathematics. ISBN: 0-89871-361-7.