

Lab 8

QR 2: Least Squares and Computing Eigenvalues

Lab Objective: *Because of its numerical stability and convenient structure, the QR decomposition is the basis of many important and practical algorithms. In this lab, we introduce linear least squares problems, tools in Python for computing least squares solutions, and two fundamental eigenvalue algorithms.*

As in the previous lab, we restrict ourselves to real matrices and therefore use the transpose in place of the Hermitian conjugate.

Least Squares

A linear system $A\mathbf{x} = \mathbf{b}$ is *overdetermined* if it has more equations than unknowns. In this situation, there is no true solution, and \mathbf{x} can only be approximated.

The *least squares solution* of $A\mathbf{x} = \mathbf{b}$, denoted as $\hat{\mathbf{x}}$, is the “closest” vector to a solution, meaning it minimizes the quantity $\|A\hat{\mathbf{x}} - \mathbf{b}\|_2$. In other words, $\hat{\mathbf{x}}$ is the vector such that $A\hat{\mathbf{x}}$ is projection of \mathbf{b} onto the range of A , and can be calculated by solving the *normal equation*:¹

$$A^T A \hat{\mathbf{x}} = A^T \mathbf{b}$$

If A is full rank, which it usually is in applications, its QR decomposition provides an efficient way to solve the normal equation. Let $A = \hat{Q}\hat{R}$ be the reduced QR decomposition of A , so \hat{Q} is $m \times n$ with orthonormal columns and \hat{R} is $n \times n$, invertible, and upper triangular. Since $\hat{Q}^T \hat{Q} = I$, and since \hat{R}^T is invertible, the normal equation can be reduced as follows (we omit the hats on \hat{Q} and \hat{R} for clarity):

$$\begin{aligned} A^T A \hat{\mathbf{x}} &= A^T \mathbf{b} \\ (QR)^T QR \hat{\mathbf{x}} &= (QR)^T \mathbf{b} \\ R^T Q^T QR \hat{\mathbf{x}} &= R^T Q^T \mathbf{b} \\ R^T R \hat{\mathbf{x}} &= R^T Q^T \mathbf{b} \\ R \hat{\mathbf{x}} &= Q^T \mathbf{b} \end{aligned} \tag{8.1}$$

¹See Volume I Chapter 3 for a formal derivation of the normal equation.

Thus $\hat{\mathbf{x}}$ is the least squares solution to $A\mathbf{x} = \mathbf{b}$ if and only if $\hat{R}\hat{\mathbf{x}} = \hat{Q}^T\mathbf{b}$. Since \hat{R} is upper triangular, this equation can be solved quickly with back substitution.

Problem 1. Write a function that accepts an $m \times n$ matrix A of rank n and a vector \mathbf{b} of length n . Use the QR decomposition and Equation 8.1 to solve the normal equation corresponding to $A\mathbf{x} = \mathbf{b}$.

You may use either SciPy's QR routine or one of your own routines from the previous lab. In addition, you may use `la.solve_triangular()`, SciPy's optimized routine for solving triangular systems.

Fitting a Line

The least squares solution can be used to find the best fit curve of a chosen type to a set of points. Consider the problem of finding the line $y = ax + b$ that best fits a set of m points $\{(x_k, y_k)\}_{k=1}^m$. Ideally, we seek a and b such that $y_k = ax_k + b$ for all k . The following linear system simultaneously represents all of these equations.

$$A\mathbf{x} = \begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ x_3 & 1 \\ \vdots & \vdots \\ x_m & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_m \end{bmatrix} = \mathbf{b} \quad (8.2)$$

Note that A has full column rank as long as not all of the x_k values are the same.

Because this system has two unknowns, it is guaranteed to have a solution if it has two or fewer equations. However, if there are more than two data points, the system is overdetermined if any set of three points are not collinear. We therefore seek a least squares solution, which in this case means finding the slope \hat{a} and y -intercept \hat{b} such that the line $y = \hat{a}x + \hat{b}$ best fits the data.

Figure 8.1 is a typical example of this idea where $\hat{a} \approx \frac{1}{2}$ and $\hat{b} \approx -3$.

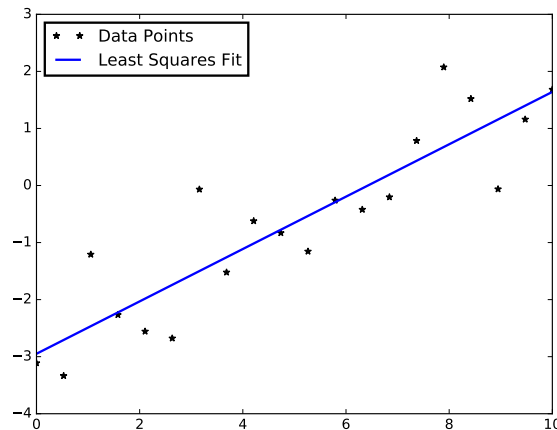


Figure 8.1

Problem 2. The file `housing.npy` contains the purchase-only housing price index, a measure of how housing prices are changing, for the United States from 2000 to 2010.^a Each row in the array is a separate measurement; the columns are the year and the price index, in that order. To avoid large numerical computations, the year measurements start at 0 instead of 2000.

Find the least squares line that relates the year to the housing price index (i.e., let year be the x -axis and index the y -axis).

1. Construct the matrix A and the vector \mathbf{b} described by Equation 8.2. (Hint: the functions `np.vstack()`, `np.column_stack()`, and/or `np.ones()` may be helpful.)
2. Use your function from Problem 1 to find the least squares solution.
3. Plot the data points as a scatter plot.
4. Plot the least squares line with the scatter plot.

^aSee <http://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx>.

NOTE

The least squares problem of fitting a line to a set of points is often called *linear regression*, and the resulting line is called the *linear regression line*. SciPy's specialized tool for linear regression is `scipy.stats.linregress()`. This function takes in an array of x -coordinates and a corresponding array of y -coordinates, and returns the slope and intercept of the regression line, along with a few other statistical measurements.

For example, the following code produces Figure 8.1.

```
>>> import numpy as np
>>> from scipy.stats import linregress

# Generate some random data close to the line y = .5x - 3.
>>> x = np.linspace(0, 10, 20)
>>> y = .5*x - 3 + np.random.randn(20)

# Use linregress() to calculate m and b, as well as the correlation
# coefficient, p-value, and standard error. See the documentation for
# details on each of these extra return values.
>>> a, b, rvalue, pvalue, stderr = linregress(x, y)

>>> plt.plot(x, y, 'k*', label="Data Points")
>>> plt.plot(x, a*x + b, 'b-', lw=2, label="Least Squares Fit")
>>> plt.legend(loc="upper left")
>>> plt.show()
```

Fitting a Polynomial

Least squares can also be used to fit a set of data to the best fit polynomial of a specified degree. Let $\{(x_k, y_k)\}_{k=1}^m$ be the set of m data points in question. The general form for a polynomial of degree n is as follows:

$$p_n(x) = c_n x^n + c_{n-1} x^{n-1} + \cdots + c_2 x^2 + c_1 x + c_0$$

Note that the polynomial is uniquely determined by its $n+1$ coefficients $\{c_k\}_{k=0}^n$. Ideally, then, we seek the set of coefficients $\{c_k\}_{k=0}^n$ such that

$$y_k = c_n x_k^n + c_{n-1} x_k^{n-1} + \cdots + c_2 x_k^2 + c_1 x_k + c_0$$

for all values of k . These m linear equations yield the following linear system:

$$A\mathbf{x} = \begin{bmatrix} x_1^n & x_1^{n-1} & \cdots & x_1^2 & x_1 & 1 \\ x_2^n & x_2^{n-1} & \cdots & x_2^2 & x_2 & 1 \\ x_3^n & x_3^{n-1} & \cdots & x_3^2 & x_3 & 1 \\ \vdots & \vdots & & \vdots & \vdots & \vdots \\ x_m^n & x_m^{n-1} & \cdots & x_m^2 & x_m & 1 \end{bmatrix} \begin{bmatrix} c_n \\ c_{n-1} \\ \vdots \\ c_2 \\ c_1 \\ c_0 \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_m \end{bmatrix} = \mathbf{b} \quad (8.3)$$

If $m > n+1$ this system is overdetermined, requiring a least squares solution.

Working with Polynomials in NumPy

The $m \times (n+1)$ matrix A of Equation 8.3 is called a *Vandermonde matrix*.² NumPy's `np.vander()` is a convenient tool for quickly constructing a Vandermonde matrix, given the values $\{x_k\}_{k=1}^m$ and the number of desired columns.

```
>>> print(np.vander([2, 3, 5], 2))
[[2 1]
 [3 1]
 [5 1]]

>>> print(np.vander([2, 3, 5, 4], 3))
[[ 4  2  1]
 [ 9  3  1]
 [25  5  1]
 [16  4  1]]
```

NumPy also has powerful tools for working efficiently with polynomials. The class `np.poly1d` represents a 1-dimensional polynomial. Instances of this class are callable like a function.³ The constructor accepts the polynomial's coefficients, from largest degree to smallest.

Table 8.1 lists the attributes and methods of the `np.poly1d` class. See <http://docs.scipy.org/doc/numpy/reference/routines.polynomials.html> for a list of NumPy's polynomial routines.

²Vandermonde matrices have many special properties and are useful for many applications, including polynomial interpolation and discrete Fourier analysis.

³Class instances can be made callable by implementing the `__call__()` magic method.

Attribute	Description
coeffs	The $n + 1$ coefficients, from greatest degree to least.
order	The polynomial degree (n).
roots	The $n - 1$ roots.
Method	Returns
deriv()	The coefficients of the polynomial after being differentiated.
integ()	The coefficients of the polynomial after being integrated (with $c_0 = 0$).

Table 8.1: Attributes and methods of the `np.poly1d` class.

```
# Create a callable object for the polynomial f(x) = (x-1)(x-2) = x^2 - 3x + 2.
>>> f = np.poly1d([1, -3, 2])
>>> print(f)
      2
1 x - 3 x + 2

# Evaluate f(x) for several values of x in a single function call.
>>> f([1, 2, 3, 4])
array([0, 0, 2, 6])

# Evaluate f(x) at 1, 2, 3, and 4 without creating f(x) explicitly.
>>> np.polyval([1, -3, 2], [1, 2, 3, 4])
array([0, 0, 2, 6])
```

Problem 3. The data in `housing.npy` is nonlinear, and might be better fit by a polynomial than a line.

Write a function that uses Equation 8.3 to calculate the polynomials of degree 3, 6, 9, and 12 that best fit the data. Plot the original data points and each least squares polynomial together in individual subplots. (Hint: define a separate, more refined domain with `np.linspace()` and use this domain to smoothly plot the polynomials.)

Instead of using Problem 1 to solve the normal equation, you may use `scipy.linalg.lstsq()`, demonstrated below.

```
>>> from scipy import linalg as la

# Define A and b appropriately.

# Solve the normal equation using SciPy's least squares routine.
# The least squares solution is the first of four return values.
>>> x = la.lstsq(A, b)[0]
```

Compare your results to `np.polyfit()`. This function receives an array of x values, an array of y values, and an integer for the polynomial degree, and returns the coefficients of the best fit polynomial of that degree.

ACHTUNG!

Having more parameters in a least squares model is not always better. For a set of m points, the best fit polynomial of degree $m - 1$ *interpolates* the data set, meaning that $p(x_k) = y_k$ exactly for each k . In this case there are enough unknowns that the system is no longer overdetermined. However, such polynomials are highly subject to numerical errors and are unlikely to accurately represent true patterns in the data.

Choosing to have too many unknowns in a fitting problem is (fittingly) called *overfitting*, and is an important issue to avoid in any statistical model.

Fitting a Circle

Suppose the set of m points $\{(x_k, y_k)\}_{k=1}^m$ are arranged in a nearly circular pattern. The general equation of a circle with radius r and center (c_1, c_2) is as follows:

$$(x - c_1)^2 + (y - c_2)^2 = r^2. \quad (8.4)$$

The circle is uniquely determined by r , c_1 , and c_2 , so these are the parameters that should be solved for in a least squares formulation of the problem. However, Equation 8.4 is not linear in any of these variables.

$$\begin{aligned} (x - c_1)^2 + (y - c_2)^2 &= r^2 \\ x^2 - 2c_1x + c_1^2 + y^2 - 2c_2y + c_2^2 &= r^2 \\ x^2 + y^2 &= 2c_1x + 2c_2y + r^2 - c_1^2 - c_2^2 \end{aligned} \quad (8.5)$$

The quadratic terms x^2 and y^2 are acceptable because the points $\{(x_k, y_k)\}_{k=1}^m$ are given. To eliminate the nonlinear terms in the unknown parameters r , c_1 , and c_2 , define a new variable $c_3 = r^2 - c_1^2 - c_2^2$. Then for each point (x_k, y_k) , Equation 8.5 becomes the following:

$$2c_1x_k + 2c_2y_k + c_3 = x_k^2 + y_k^2$$

These m equations are linear in c_1 , c_2 , and c_3 , and can be written as a linear system.

$$\begin{bmatrix} 2x_1 & 2y_1 & 1 \\ 2x_2 & 2y_2 & 1 \\ \vdots & \vdots & \vdots \\ 2x_m & 2y_m & 1 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} x_1^2 + y_1^2 \\ x_2^2 + y_2^2 \\ \vdots \\ x_m^2 + y_m^2 \end{bmatrix} \quad (8.6)$$

After solving for the least squares solution, r can be recovered with the relation $r = \sqrt{c_1^2 + c_2^2 + c_3}$. Finally, plotting a circle is best done with polar coordinates. Using the same variables as before, the circle can be represented in polar equations with the following equations:

$$x = r \cos(\theta) + c_1 \quad y = r \sin(\theta) + c_2,$$

where $\theta \in [0, 2\pi]$.

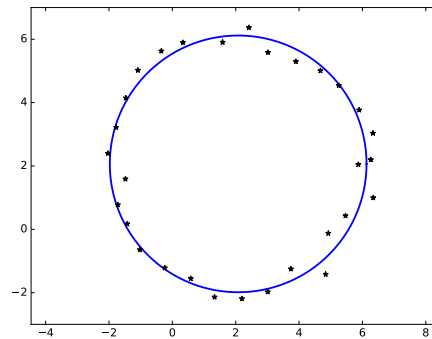
```

# Load some data and construct the matrix A and the vector b.
>>> xk, yk = np.load("circle.npy").T
>>> A = np.column_stack((2*x, 2*y, np.ones_like(x)))
>>> b = xk**2 + yk**2

# Calculate the least squares solution and calculate the radius.
>>> c1, c2, c3 = la.lstsq(A, b)[0]
>>> r = np.sqrt(c1**2 + c2**2 + c3)

# Plot the circle using polar coordinates.
>>> theta = np.linspace(0, 2*np.pi, 200)
>>> x = r*np.cos(theta) + c1
>>> y = r*np.sin(theta) + c2
>>> plt.plot(x, y, '-', lw=2)
>>> plt.plot(xk, yk, 'k*')
>>> plt.axis("equal")

```



Problem 4. The general equation for an ellipse is

$$ax^2 + bx + cxy + dy + ey^2 = 1.$$

Write a function that calculates the parameters for the ellipse that best fits the data in the file `ellipse.npy`. Plot the original data points and the ellipse together, using the following function to plot the ellipse.

```

def plot_ellipse(a, b, c, d, e):
    """Plot an ellipse of the form ax^2 + bx + cxy + dy + ey^2 = 1."""
    theta = np.linspace(0, 2*np.pi, 200)
    cos_t, sin_t = np.cos(theta), np.sin(theta)
    A = a*(cos_t**2) + c*cos_t*sin_t + e*(sin_t**2)
    B = b*cos_t + d*sin_t
    r = (-B + np.sqrt(B**2 + 4*A))/(2*A)

    plt.plot(r*cos_t, r*sin_t, lw=2)
    plt.gca().set_aspect("equal", "datalim")

```

Computing Eigenvalues

The eigenvalues of an $n \times n$ matrix A are the roots of its characteristic polynomial $\det(A - \lambda I)$. Thus, finding the eigenvalues of A amounts to computing the roots of a polynomial of degree n . However, for $n \geq 5$, it is provably impossible to find an algebraic closed-form solution to this problem.⁴ In addition, numerically computing the roots of a polynomial is a famously ill-conditioned problem, meaning that small changes in the coefficients of the polynomial (brought about by small changes in the entries of A) may yield wildly different results. Instead, eigenvalues must be computed with iterative methods.

The Power Method

The *dominant eigenvalue* of the $n \times n$ matrix A is the unique eigenvalue of greatest magnitude, if such an eigenvalue exists. The *power method* iteratively computes the dominant eigenvalue of A and its corresponding eigenvector.

Begin by choosing a vector \mathbf{x}_0 such that $\|\mathbf{x}_0\| = 1$, and define the following:

$$\mathbf{x}_{k+1} = \frac{A\mathbf{x}_k}{\|A\mathbf{x}_k\|}$$

If A has a dominant eigenvalue λ , and if the projection of \mathbf{x}_0 onto the subspace spanned by the eigenvectors corresponding to λ is nonzero, then the sequence of vectors $\{\mathbf{x}_k\}_{k=0}^{\infty}$ converges to an eigenvector \mathbf{x} of A corresponding to λ .

Since \mathbf{x} is an eigenvector of A , $A\mathbf{x} = \lambda\mathbf{x}$. Left multiplying by \mathbf{x}^T on each side gives $\mathbf{x}^T A\mathbf{x} = \lambda\mathbf{x}^T \mathbf{x}$, and hence $\lambda = \frac{\mathbf{x}^T A\mathbf{x}}{\mathbf{x}^T \mathbf{x}}$. This ratio is called the *Rayleigh quotient*. However, since each \mathbf{x}_k is normalized, $\mathbf{x}^T \mathbf{x} = \|\mathbf{x}\|^2 = 1$, so $\lambda = \mathbf{x}^T A\mathbf{x}$.

The entire algorithm is summarized below.

Algorithm 8.1

```

1: procedure POWER METHOD( $A$ )
2:    $m, n \leftarrow \text{shape}(A)$                                  $\triangleright A$  is square so  $m = n$ .
3:    $\mathbf{x}_0 \leftarrow \text{random}(n)$                                 $\triangleright A$  random vector of length  $n$ 
4:    $\mathbf{x}_0 \leftarrow \mathbf{x}_0 / \|\mathbf{x}_0\|$                                 $\triangleright$  Normalize  $\mathbf{x}_0$ 
5:   for  $k = 1, 2, \dots, N - 1$  do
6:      $\mathbf{x}_{k+1} \leftarrow A\mathbf{x}_k$ 
7:      $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_{k+1} / \|\mathbf{x}_{k+1}\|$ 
8:   return  $\mathbf{x}_N^T A\mathbf{x}_N, \mathbf{x}_N$ 

```

The power method is limited by a few assumptions. First, not all square matrices A have a dominant eigenvalue. However, the Perron-Frobenius theorem guarantees that if all entries of A are positive, then A has a dominant eigenvalue. Second, there is no way to choose an \mathbf{x}_0 that is guaranteed to have a nonzero projection onto the span of the eigenvectors corresponding to λ , though a random \mathbf{x}_0 will almost surely satisfy this condition. Even with these assumptions, a rigorous proof that the power method converges is most convenient with tools from spectral calculus. See Chapter ?? of the Volume I text for details.

⁴This result, called *Abel's impossibility theorem*, was first proven by Niels Henrik Abel in 1824.

Problem 5. Write a function that accepts an $n \times n$ matrix A , a maximum number of iterations N , and a stopping tolerance tol . Use Algorithm 8.1 to compute the dominant eigenvalue of A and a corresponding eigenvector.

Continue the loop in step 5 until either $\|\mathbf{x}_{k+1} - \mathbf{x}_k\|$ is less than the tolerance tol , or until iterating the maximum number of times N .

Test your function on square matrices with all positive entries, verifying that $A\mathbf{x} = \lambda\mathbf{x}$. Use SciPy's eigenvalue solver, `scipy.linalg.eig()`, to compute all of the eigenvalues and corresponding eigenvectors of A and check that λ is the dominant eigenvalue of A .

```
# Construct a random matrix with positive entries.
>>> A = np.random.random((10,10))

# Compute the eigenvalues and eigenvectors of A via SciPy.
>>> eigs, vecs = la.eig(A)

# Get the dominant eigenvalue and eigenvector of A.
# The eigenvector of the kth eigenvalue is the kth column of 'vecs'.
>>> loc = np.argmax(eigs)
>>> lamb, x = eigs[loc], vecs[:,loc]

# Verify that Ax = lambda x.
>>> np.allclose(A.dot(x), lamb*x)
True
```

The QR Algorithm

An obvious shortcoming of the power method is that it only computes the largest eigenvalue and a corresponding eigenvector. The QR algorithm, on the other hand, attempts to find all eigenvalues of A .

Let $A_0 = A$, and for arbitrary k let $Q_k R_k = A_k$ be the QR decomposition of A_k . Since A is square, so are Q_k and R_k , so they can be recombined in reverse order.

$$A_{k+1} = R_k Q_k$$

This recursive definition establishes an important relation between the A_k .

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

Thus A_k is orthonormally similar to A_{k+1} , and similar matrices have the same eigenvalues. The series of matrices $\{A_k\}_{k=0}^{\infty}$ converges to the following block matrix.

$$S = \begin{bmatrix} S_1 & * & \cdots & * \\ \mathbf{0} & S_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & * \\ \mathbf{0} & \cdots & \mathbf{0} & S_m \end{bmatrix}$$

Each S_i is either a 1×1 or 2×2 matrix.⁵ Since S is block upper triangular, its eigenvalues are the eigenvalues of its diagonal S_i blocks. Then because A is similar to each A_k , those eigenvalues of S are the eigenvalues of A .

When A has real entries but complex eigenvalues, 2×2 S_i blocks appear in S . Finding eigenvalues of a 2×2 matrix is equivalent to finding the roots of a 2nd degree polynomial, which has a closed form solution via the quadratic equation. This implies that complex eigenvalues come in conjugate pairs.

$$\begin{aligned} \det(S_i - \lambda I) &= \begin{vmatrix} a - \lambda & b \\ c & d - \lambda \end{vmatrix} = (a - \lambda)(d - \lambda) - bc \\ &= \lambda^2 - ad\lambda + (ad - bc) \end{aligned} \quad (8.7)$$

Hessenberg Preconditioning

The QR algorithm works more accurately and efficiently on matrices that are in upper Hessenberg form, as upper Hessenberg matrices are already close to triangular. Furthermore, if $H = QR$ is the QR decomposition of upper Hessenberg H then RQ is also upper Hessenberg, so the almost-triangular form is preserved at each iteration. Putting a matrix in upper Hessenberg form before applying the QR algorithm is called *Hessenberg preconditioning*.

With preconditioning in mind, the entire QR algorithm is as follows.

Algorithm 8.2

```

1: procedure QR_ALGORITHM( $A$ )
2:    $m, n \leftarrow \text{shape}(A)$ 
3:    $S \leftarrow \text{hessenberg}(A)$  ▷ Put  $A$  in upper Hessenberg form.
4:   for  $k = 0, 1, \dots, N - 1$  do
5:      $Q, R \leftarrow \text{qr}(S)$  ▷ Get the QR decomposition of  $A_k$ .
6:      $S \leftarrow RQ$  ▷ Recombine  $R_k$  and  $Q_k$  into  $A_{k+1}$ .
7:    $\text{eigs} \leftarrow []$  ▷ Initialize an empty list of eigenvalues.
8:    $i \leftarrow 0$ 
9:   while  $i < n$  do
10:    if  $S_i$  is  $1 \times 1$  then
11:      Append  $S_i$  to  $\text{eigs}$ 
12:    else if  $S_i$  is  $2 \times 2$  then
13:      Calculate the eigenvalues of  $S_i$ 
14:      Append the eigenvalues of  $S_i$  to  $\text{eigs}$ 
15:       $i \leftarrow i + 1$ 
16:       $i \leftarrow i + 1$  ▷ Move to the next  $S_i$ .
17:   return  $\text{eigs}$ 

```

⁵If all of the S_i are 1×1 matrices, then the upper triangular S is called the *Schur form* of A . If some of the S_i are 2×2 matrices, then S is called the *real Schur form* of A .

Problem 6. Write a function that accepts an $n \times n$ matrix A , a number of iterations N , and a tolerance tol . Use Algorithm 8.2 to implement the QR algorithm with Hessenberg preconditioning, returning the eigenvalues of A .

Consider the following implementation details.

- Use `scipy.linalg.hessenberg()` to reduce A to upper Hessenberg form in step 3, or use your own Hessenberg function from the previous lab.
- The loop in step 4 should run for N total iterations.
- Use `scipy.linalg.qr()` to compute the QR decomposition of S in step 5, or use your own QR factorization routine from the previous lab (since S is upper Hessenberg, Givens rotations are the most efficient way to produce Q and R).
- Assume that S_i is 1×1 in step 10 if one of two criteria hold:
 1. S_i is the last diagonal entry of S .
 2. The absolute value of element below the i th main diagonal entry of S (the lower left element of the 2×2 block) is less than tol .
- If S_i is 2×2 , use the quadratic formula and Equation 8.7 to compute its eigenvalues. Use the function `cmath.sqrt()` to correctly compute the square root of a negative number.

Test your function on small random symmetric matrices, comparing your results to `scipy.linalg.eig()`. To construct a random symmetric matrix, note that $A + A^T$ is always symmetric.

NOTE

Algorithm 8.2 is theoretically sound, but can still be greatly improved. Most modern computer packages instead use the *implicit QR algorithm*, an improved version of the QR algorithm, to compute eigenvalues.

For large matrices, there are other iterative methods besides the power method and the QR algorithm for efficiently computing eigenvalues. They include the Arnoldi iteration, the Jacobi method, the Rayleigh quotient method, and others. We will return to this subject after studying spectral calculus and Krylov subspaces.