

Lab 1

GMRES

Lab Objective: *In this lab we will learn how to use the GMRES algorithm.*

The GMRES ("Generalized Minimal Residuals") algorithm is an efficient way to solve large linear systems. It is an iterative method that uses Krylov subspaces to reduce a high-dimensional problem to a sequence of smaller dimensional problems.

The GMRES Algorithm

Let A be an invertible $m \times m$ matrix and let \mathbf{b} be an m -vector. Let $\mathcal{K}_n(A, \mathbf{b})$ be the order- n Krylov subspace generated by A and \mathbf{b} . The idea of the GMRES algorithm is that instead of solving $A\mathbf{x} = \mathbf{b}$ directly, we use least squares to find $\mathbf{x}_n \in \mathcal{K}_n$ that minimizes the residual $r_n = \|\mathbf{b} - A\mathbf{x}_n\|_2$. The algorithm returns when this residual is sufficiently small. In good circumstances, this will happen when n is still much less than m .

The GMRES algorithm is implemented with the Arnoldi iteration for numerical stability. The Arnoldi iteration produces H_n , an $(n+1) \times n$ upper Hessenberg matrix, and Q_n , the matrix containing the basis vectors of $\mathcal{K}_n(A, \mathbf{b})$, such that $AQ_n = Q_{n+1}H_n$. We are looking for $\mathbf{x}_n = Q_n\mathbf{y}_n + \mathbf{x}_0$ for some $\mathbf{y}_n \in \mathbb{R}^n$ which minimizes the norm of $\mathbf{b} - A\mathbf{x}_n$. Since the columns of Q are orthonormal, we can compute the residual equivalently as

$$\|\mathbf{b} - A\mathbf{x}_n\|_2 = \|Q_{n+1}(\beta\mathbf{e}_1 - H_n\mathbf{y}_n)\|_2 = \|H_n\mathbf{y}_n - \beta\mathbf{e}_1\|_2. \quad (1.1)$$

Here \mathbf{e}_1 is the vector $(1, 0, \dots, 0)$ of length $n+1$. β is the Euclidean norm of $\mathbf{b} - A\mathbf{x}_0$, where \mathbf{x}_0 is an initial arbitrary guess of the solution. (Ordinarily this guess is zero, and then the $A\mathbf{x}_0$ could be left out; however, a modified version of the algorithm will be discussed at the end of the lab, in which other nonzero guesses will be made.) Thus to minimize the left side of 1.1, we can minimize the right, and \mathbf{x}_n can be computed as $Q_n\mathbf{y}_n + \mathbf{x}_0$.

This algorithm is outlined in Algorithm 1.1. For a complete derivation see [TODO: ref textbook].

Algorithm 1.1 The GMRES algorithm. This algorithm operates on a vector \mathbf{b} and matrix A . It iterates k times or until the residual is less than tol , returning an approximate solution to $A\mathbf{x} = \mathbf{b}$ and the error in this approximation.

```

1: procedure GMRES( $A, \mathbf{b}, \mathbf{x}_0, k, tol$ )
2:    $Q \leftarrow \text{empty}(\text{size}(\mathbf{b}), k + 1)$  ▷ Initialize
3:    $H \leftarrow \text{zeros}(k + 1, k)$ 
4:    $\mathbf{r}_0 \leftarrow \mathbf{b} - A\mathbf{x}_0$ 
5:    $Q[:, 0] = \mathbf{r}_0 / \|\mathbf{r}_0\|_2$ 
6:   for  $n = 1 \dots k$  do
7:     Set entries of  $Q$  and  $H$  as in Arnoldi iteration.
8:     Compute the residual  $res$  and the least squares solution  $\mathbf{y}_n$  for the part
       of  $H$  so far created (equation 1.1).
9:     if  $res < tol$  then
10:      return  $Q[:, : n + 1]\mathbf{y} + \mathbf{x}_0, res$ 
11:  return  $Q[:, : n + 1]\mathbf{y} + \mathbf{x}_0, res$ 

```

Problem 1. Use Algorithm 1.1 to complete the following Python function implementing the GMRES algorithm.

```

def gmres(A, b, x0, k=100, tol=1e-8):
    '''Calculate approximate solution of Ax=b using GMRES algorithm.

    INPUTS:
    A      - Callable function that calculates Ax for any input vector x.
    b      - A NumPy array of length m.
    x0     - An arbitrary initial guess.
    k      - Maximum number of iterations of the GMRES algorithm. Defaults ↔
              to 100.
    tol    - Stop iterating if the residual is less than 'tol'. Defaults to ↔
              1e-8.

    RETURN:
    Return (y, res) where 'y' is an approximate solution to Ax=b and 'res'
    is the residual.

    Examples:
    >>> a = np.array([[1,0,0],[0,2,0],[0,0,3]])
    >>> A = lambda x: a.dot(x)
    >>> b = np.array([1, 4, 6])
    >>> x0 = np.zeros(b.size)
    >>> gmres(A, b, x0)
    (array([ 1.,  2.,  2.]), 1.09808907533e-16)
    '''

```

You may assume that the input \mathbf{b} is a real array and the function $A()$ always outputs real arrays.

Hint: Use `numpy.linalg.lstsq()` to solve the least squares problem. Be sure to read the documentation so you know what the function returns to you.

Convergence of GMRES

At the n -th iteration, GMRES computes the best approximate solution $\mathbf{x} \in \mathcal{K}_n$ to $A\mathbf{x} = \mathbf{b}$. If A is full rank, then $\mathcal{K}_m = \mathbb{F}^m$, so the m^{th} iteration will always return an exact answer. However, we say the algorithm converges after n steps if the n^{th} residual is sufficiently small.

The rate of convergence of GMRES depends on the eigenvalues of A .

Problem 2. Implement the following Python function by modifying your solution to Problem 1.

```
def plot_gmres(A, b, x0, tol=1e-8):
    '''Use the GMRES algorithm to approximate the solution to Ax=b. Plot ↵
        the eigenvalues of A and the convergence of the algorithm.

    INPUTS:
    A   - A 2-D NumPy array of shape mxm.
    b   - A 1-D NumPy array of length m.
    x0  - An arbitrary initial guess.
    tol - Stop iterating and create the desired plots when the residual is
          less than 'tol'. Defaults to 1e-8.

    OUTPUT:
    Follow the GMRES algorithm until the residual is less than tol, for a
    maximum of m iterations. Then create the two following plots (subplots
    of a single figure):

    1. Plot the eigenvalues of A in the complex plane.

    2. Plot the convergence of the GMRES algorithm by plotting the
       iteration number on the x-axis and the residual on the y-axis.
       Use a log scale on the y-axis.
    '''
```

Use this function to investigate the convergence of GMRES as follows. Define an $m \times m$ matrix

$$A_n = nI + P,$$

where I is the $m \times m$ identity matrix and P is a $m \times m$ matrix of numbers from a random normal distribution with mean 0 and standard deviation $1/(2\sqrt{m})$. Write a function that calls `plot_gmres` on A_n for $n = -4, -2, 0, 2, 4$. Use $m = 200$, let \mathbf{b} be an array of ones, and let \mathbf{x}_0 be the zero vector or anything else that suits you. How does the convergence of the GMRES algorithm relate to the eigenvalues?

Hints:

1. Create a plot with a log scale on the y-axis with `plt.yscale('log')`.
2. Create a matrix with entries from a random normal distribution with `np.random.normal()`. Read the documentation for more information.

3. Note that the parameter A required here is not a callable function but a matrix; this is to allow the finding of the eigenvalues.
4. Output for $n = 2$, $m = 200$ is in Figure 1.1 below.

Ideas for this problem were taken from Example 35.1 on p. 271 of [Trefethen1997].




Figure 1.1: The left plot is the eigenvalues of the matrix A_2 , which is defined in Problem 2. The right plot is the convergence of the GMRES algorithm on A_2 with starting vector $\mathbf{b} = (1, 1, \dots, 1)$. This figure is one possible output of the function `plot_gmres()`.

Improving GMRES

There are many ways to make the GMRES algorithm more robust and efficient.

Breakdowns in GMRES

One of the selling points of GMRES is that it can't break down unless it reaches an exact solution. In other words, the only way GMRES could break down is if a vector found by the Arnoldi iteration is 0. That is, suppose, we have already computed

$$\mathcal{K}_n(A, \mathbf{b}) = \text{span}\{\mathbf{b}, A\mathbf{b}, \dots, A^{n-1}\mathbf{b}\} = \text{span}\{\mathbf{q}_1, \dots, \mathbf{q}_n\}.$$

We next compute $A^n\mathbf{b}$ and orthogonalize it against $\mathcal{K}_n(A, \mathbf{b})$, yielding \mathbf{q}_{n+1} . But if $A^n\mathbf{b} \in \mathcal{K}_n(A, \mathbf{b})$ then \mathbf{q}_{n+1} will be 0, and our algorithm will break when we try to normalize \mathbf{q}_{n+1} .

In this situation, the least squares solution to (1.1) is an *exact* solution to $A\mathbf{x} = \mathbf{b}$. In other words, \mathbf{b} is in the $\text{span}\{\mathbf{q}_1, \dots, \mathbf{q}_n\}$. Fortunately, precautions against this have been taken in our implementation of the Arnoldi algorithm.

GMRES with Restarts

The first few iterations of GMRES have low spatial and temporal complexity. However, as k increases, the k^{th} iteration of GMRES becomes more expensive in both time and memory. In fact, computing the k^{th} iteration of GMRES for very large k can be prohibitively complex.

This issue is addressed by using GMRES(k), or GMRES with restarts. When k becomes large, this algorithm restarts GMRES but with an improved initial guess. GMRES with restarts is outlined in Algorithm 1.2.

Algorithm 1.2 The GMRES(k) algorithm. This algorithm performs GMRES on a vector \mathbf{b} and matrix A . It iterates k times before restarting. It terminates after *restarts* restarts or when the residual is less than *tol*, returning an approximate solution to $A\mathbf{x} = \mathbf{b}$ and the error in this approximation.

```

1: procedure GMRES( $\kappa$ )( $A, \mathbf{b}, \mathbf{x}_0, k, \text{tol}, \text{restarts}$ )
2:    $n \leftarrow 0$  ▷ Initialize
3:   while  $n \leq \text{restarts}$  do
4:     Perform the GMRES algorithm, obtaining a least squares solution  $\mathbf{y}$ .
5:     If the desired tolerance was reached, return. Otherwise, continue.
6:      $\mathbf{x}_0 \leftarrow \mathbf{y}$ 
7:      $n \leftarrow n + 1$ 
8:   return  $\mathbf{y}, \text{res}$  ▷ Return the approximate solution and the residual

```

The algorithm GMRES(k) will always have manageable spatial and temporal complexity, but it is less reliable than GMRES. If the true solution \mathbf{x} to $A\mathbf{x} = \mathbf{b}$ is nearly orthogonal to the Krylov subspaces $\mathcal{K}_n(A, \mathbf{b})$ for $n \leq k$, then GMRES(k) could converge very slowly or not at all.

Problem 3. Implement Algorithm 1.2 with the following function.

```

def gmres_k(A, b, x0, k=5, tol=1E-8, restarts=50):
    '''Use the GMRES(k) algorithm to approximate the solution to Ax=b.

    INPUTS:
    A      - A callable function that calculates Ax for any vector x.
    b      - A NumPy array.
    x0     - An arbitrary initial guess.
    k      - Maximum number of iterations of the GMRES algorithm before
              restarting. Defaults to 5.
    tol    - Stop iterating if the residual is less than 'tol'. Defaults
              to 1E-8.
    restarts - Maximum number of restarts. Defaults to 50.

    RETURN:
    Return (y, res) where 'y' is an approximate solution to Ax=b and 'res'

```

```
is the residual.
'''
```

Compare the speed of `gmres()` from Problem 1 and `gmres_k()` on the matrices in Problem 2.

GMRES in SciPy

The GMRES algorithm is implemented in SciPy as the function `scipy.sparse.linalg.gmres()`. Here we use this function to solve $A\mathbf{x} = \mathbf{b}$ where A is a random 300×300 matrix and \mathbf{b} is a random vector.

```
>>> import numpy as np
>>> from scipy import sparse as spar
>>> from scipy import linalg as la
>>>
>>> A = np.random.rand(300, 300)
>>> b = np.random(300)
>>> x, info = spar.linalg.gmres(A, b)
>>> info
3000
```

The function outputs two objects: the approximate solution \mathbf{x} and a constant `info` telling if the function converged. If `info=0` then convergence occurred; if `info` is positive then it equals the number of iterations performed. In this case, the function performed 3000 iterations of GMRES before returning the approximate solution \mathbf{x} . We can check how close the solution is.

```
>>> la.norm(A.dot(x)-b)
4.744196381683801
```

We can get a better approximation using GMRES with restarts.

```
>>> # Restart after 1000 iterations
>>> x, info = spar.linalg.gmres(A, b, restart=1000)
>>> info
0
>>> la.norm(A.dot(x)-b)
1.0280404494143551e-12
```

This time, the returned approximation \mathbf{x} is about as close to a true solution as we could hope for.