3. Solving linear systems

I. Matrix concepts

Matrix

Defined as a table of numbers, arranged in rows and columns.

- Plural: matrices
- Conventionally denoted by bold (or double underlined in handwriting) capital letters

Example:

$$\mathbf{M} = \begin{pmatrix} 1 & -1 & 2 & 2 \\ 1 & 0.3 & 0.1 & 1 \\ 1 & 1 & 0 & -1 \end{pmatrix}$$

An element in a matrix can be denoted by the row and column indices. For example, $M_{2,3} = 0.1$

A matrix is *square* if the number of columns is equal to the number of rows.

The *main diagonal* of a matrix includes the elements whose row index is the same as their column index.

Diagonal matrix

This is a matrix where all elements not in the main diagonal are zero.
Example:

$$\begin{pmatrix} 9 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 10 & 0 \\ 0 & 0 & 0 & 8 \end{pmatrix}$$

Identity matrix

This is a diagonal square matrix where all elements on the main diagonal are 1.
Example:

$$\mathbf{I}_4 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Upper triangular matrix

This is a matrix where everything in the lower triangle (below main diagonal, row index greater than column index) is 0.
Example:

$$\begin{pmatrix} 9 & 6 & 2 & 6 \\ 0 & 8 & 3 & 2 \\ 0 & 0 & 3 & 2 \\ 0 & 0 & 0 & 8 \end{pmatrix}$$

Lower triangular matrix

This is a matrix where everything in the upper triangle (above main diagonal, row index less than column index) is 0.
Example:

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 3 & -4 & 0 & 0 \\ -2 & 0 & 2 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

Transpose

The transpose of a matrix is formed by interchanging its row and column indices.
Example:

$$\text{If } \mathbf{M} = \begin{pmatrix} 1 & -1 & 2 & 2 \\ 1 & 0.3 & 0.1 & 1 \\ 1 & 1 & 0 & -1 \end{pmatrix}, \text{then } \mathbf{M}^T = \begin{pmatrix} 1 & 1 & 1 \\ -1 & 0.3 & 1 \\ 2 & 0.1 & 0 \\ 2 & 1 & -1 \end{pmatrix}$$

A *row vector* is a matrix with only one row. Its transpose would be a *column vector*, with only one column.

A *symmetric* matrix is one that is the same as its transpose.
Example:

$$\begin{pmatrix} 9 & 3 & 1 & 3 \\ 3 & 5 & 2 & 8 \\ 1 & 2 & 6 & 4 \\ 3 & 8 & 4 & 7 \end{pmatrix}$$

Matrix addition

We can add a number (scalar) to each element in a matrix:

$$3 + \begin{pmatrix} 1 & -1 & 2 & 2 \\ 1 & 0.3 & 0.1 & 1 \\ 1 & 1 & 0 & -1 \end{pmatrix} = \begin{pmatrix} 4 & 2 & 5 & 5 \\ 4 & 3.3 & 3.1 & 4 \\ 4 & 4 & 3 & 2 \end{pmatrix}$$

Or we can add corresponding elements in two matrices that are the same *size* (number of rows and columns)

$$\begin{pmatrix} -1 & -1 & 0 & 2 \\ -1 & -1 & 2 & 3 \\ -1 & 0 & 1 & 0 \end{pmatrix} + \begin{pmatrix} 1 & -1 & 2 & 2 \\ 1 & 0.3 & 0.1 & 1 \\ 1 & 1 & 0 & -1 \end{pmatrix} = \begin{pmatrix} 0 & -2 & 2 & 4 \\ 0 & -0.7 & 2.1 & 4 \\ 0 & 1 & 1 & -1 \end{pmatrix}$$

Similarly for subtraction.

Scalar and elementwise multiplication

We can multiply each element of a matrix by a number:

$$2\begin{pmatrix} 1 & -1 & 2 & 2 \\ 1 & 0.3 & 0.1 & 1 \\ 1 & 1 & 0 & -1 \end{pmatrix} = \begin{pmatrix} 2 & -2 & 4 & 4 \\ 2 & 0.6 & 0.2 & 2 \\ 2 & 2 & 0 & -2 \end{pmatrix}$$

Or we can multiply corresponding elements in two matrices that are the same size:

$$\begin{pmatrix} -1 & -1 & 0 & 2 \\ -1 & -1 & 2 & 3 \\ -1 & 0 & 1 & 0 \end{pmatrix} \odot \begin{pmatrix} 1 & -1 & 2 & 2 \\ 1 & 0.3 & 0.1 & 1 \\ 1 & 1 & 0 & -1 \end{pmatrix} = \begin{pmatrix} -1 & 1 & 0 & 4 \\ -1 & -0.3 & 0.2 & 3 \\ -1 & 0 & 0 & 0 \end{pmatrix}$$

Similarly for division.

Matrix multiplication

This is different from elementwise multiplication. Each element in the result is the dot product of the corresponding row of the first matrix with the corresponding column of the second matrix, notated as $(AB)_{i,j} = \sum_{k=1}^{p} A_{i,k} B_{k,j}$.
p the number of columns in $\mathbf{A}$ and the number of rows in $\mathbf{B}$, which must be the same.
For example,

$$\begin{pmatrix} -1 & -1 & 0 & 2 \\ -1 & -1 & 2 & 3 \\ -1 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \\ -1 & 0.3 & 1 \\ 2 & 0.1 & 0 \\ 2 & 1 & -1 \end{pmatrix} = \begin{pmatrix} 4 & 0.7 & -4 \\ 10 & 1.9 & -5 \\ 1 & -0.9 & -1 \end{pmatrix}.$$

Matrix multiplication is *associative*, i.e. $(\mathbf{AB})\mathbf{C} = \mathbf{A}(\mathbf{BC})$, but not *commutative*, i.e. $\mathbf{BA} \neq \mathbf{AB}$
If $\mathbf{A}$ has n columns and $\mathbf{I}_n$ is the $n \times n$ identity matrix, $\mathbf{AI}_n = \mathbf{A}$.
Similarly, if $\mathbf{A}$ has n rows, $\mathbf{I}_n\mathbf{A} = \mathbf{A}$

Matrix inverse

An $n \times n$ matrix $\mathbf{B}$ is said to be the inverse of a given matrix $\mathbf{A}$ of the same size if $\mathbf{AB} = \mathbf{BA} = \mathbf{I}_n$

For example,

$$\begin{pmatrix} -1 & -2 & 1 \\ -1 & 0 & 1 \\ -4 & 2 & -1 \end{pmatrix} \begin{pmatrix} -0.2 & 0 & -0.2 \\ -0.5 & 0.5 & 0 \\ -0.2 & 1 & -0.2 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

"Most" square matrices have an inverse. A matrix can only have one inverse.

Sparse matrices

A matrix is said to be sparse if most of its elements are zero.

Sparsity is relevant to computation because if a matrix has many elements but the vast majority are zero, storing only the nonzero elements (and their row and column indices) can save memory.

Tridiagonal matrices

A tridiagonal matrix only has nonzero entries on the main diagonal and directly above and below it (i.e., where the row index and column index differ by no more than 1).

For example,

$$\begin{pmatrix} 4 & 1 & 0 & 0 & 0 \\ 7 & 8 & 5 & 0 & 0 \\ 0 & 5 & 9 & 8 & 0 \\ 0 & 0 & 2 & 8 & 1 \\ 0 & 0 & 0 & 5 & 7 \end{pmatrix}$$

II. Solving linear systems

Systems

Typically in an engineering problem, there are many quantities that all play a part. Instead of just changing one at a time, we need to think of a system that includes them all.

If these quantities are denoted $x_1, x_2, x_3, \ldots, x_n$, the simplest mathematical model for a system includes linear combinations of them, of the form
$a_1x_1 + a_2x_2 + a_3x_3 + \ldots + a_nx_n = b$
where the a_i and b are some known values (a linear equation).

System of linear equations

If we have n quantities and n different linear equations involving them, there is typically a unique solution to the system. We can write the system generically as

$$\begin{aligned}
a_{1,1}x_1 + a_{2,1}x_2 + a_{3,1}x_3 + \ldots + a_{n,1}x_n &= b_1 \\
a_{1,2}x_1 + a_{2,2}x_2 + a_{3,2}x_3 + \ldots + a_{n,2}x_n &= b_2 \\
a_{1,3}x_1 + a_{2,3}x_2 + a_{3,3}x_3 + \ldots + a_{n,3}x_n &= b_3 \\
&\ldots \\
a_{1,n}x_1 + a_{2,n}x_2 + a_{3,n}x_3 + \ldots + a_{n,n}x_n &= b_n,
\end{aligned}$$

where for each equation j, the $a_{i,j}$ coefficients are different.

Using the definition of matrix multiplication, we can write this more compactly as $\mathbf{Ax} = \mathbf{b}$, where $\mathbf{A}$ is the $n \times n$ matrix whose row i, column j element is $a_{i,j}$, $\mathbf{x}$ is the $n \times 1$ matrix of problem quantities whose ith element is x_i, and $\mathbf{b}$ is the $n \times 1$ matrix whose jth element is b_j.

Example

If the linear equations are

$$\begin{aligned} 2x_1 + 3x_3 &= 17 \\ -x_1 + x_2 - 2x_3 &= -4 \\ x_1 + 2x_2 - 4x_3 &= -5, \end{aligned}$$

we can write them as a matrix multiplication:

$$\begin{pmatrix} 2 & 0 & 3 \\ -1 & 1 & -2 \\ 1 & 2 & -4 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 17 \\ -4 \\ -5 \end{pmatrix}$$

Or, more compactly, we can not show the unknown x_i, giving the *augmented matrix* form

$$\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ -1 & 1 & -2 & -4 \\ 1 & 2 & -4 & -5 \end{array}\right)$$

Solving linear systems: diagonal coefficient matrix

If $\mathbf{Ax} = \mathbf{b}$ and $\mathbf{A}$ is a diagonal matrix, we can solve for $\mathbf{x}$ quickly with just n divisions of the form $x_i = b_i / A_{i,i}$. There is a unique solution as long as none of the diagonal elements $A_{i,i}$ is zero. For example,

$$\begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -4 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 17 \\ -4 \\ -5 \end{pmatrix}$$

Solving linear systems: upper triangular coefficient matrix

If $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\mathbf{A}$ is an upper triangular matrix, we can find $\mathbf{x}$ with a total of n^2 arithmetic operations as follows (*back substitution*):

$$
\begin{aligned}
x_n &= b_n / A_{n,n} \\
x_{n-1} &= (b_{n-1} - A_{n-1,n} x_n) / A_{n-1,n-1} \\
x_{n-2} &= (b_{n-2} - A_{n-2,n} x_n - A_{n-2,n-1} x_{n-1}) / A_{n-2,n-2} \\
\dots \quad & [x_i = (b_i - \sum_{j=i+1}^{n} A_{i,j} x_j) / A_{i,i}] \\
x_1 &= (b_1 - \sum_{j=2}^{n} A_{1,j} x_j) / A_{1,1}
\end{aligned}
$$

There is a unique solution as long as none of the diagonal elements $A_{i,i}$ is zero.

For example,

$$
\begin{pmatrix} 2 & 0 & 3 \\ 0 & 1 & -2 \\ 0 & 0 & -4 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 17 \\ -4 \\ -5 \end{pmatrix}
$$

Solving linear systems: lower triangular coefficient matrix

If $\mathbf{Ax} = \mathbf{b}$ and $\mathbf{A}$ is a lower triangular matrix, we can find $\mathbf{x}$ with a total of n^2 arithmetic operations as follows (*forward substitution*):

$$\begin{aligned}
x_1 &= b_1/A_{1,1} \\
x_2 &= (b_2 - A_{1,2}x_1)/A_{2,2} \\
x_3 &= (b_3 - A_{1,3}x_1 - A_{2,3}x_2)/A_{3,3} \\
\ldots \quad & [x_i = (b_i - \sum_{j=1}^{i-1} A_{i,j}x_j)/A_{i,i}] \\
x_n &= (b_n - \sum_{j=1}^{n-1} A_{n,j}x_j)/A_{n,n}
\end{aligned}$$

There is a unique solution as long as none of the diagonal elements $A_{i,i}$ is zero.

For example,

$$\begin{pmatrix} 2 & 0 & 0 \\ -1 & 1 & 0 \\ 1 & 2 & -4 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 17 \\ -4 \\ -5 \end{pmatrix}$$

Solving general linear systems

If $\mathbf{Ax} = \mathbf{b}$ and $\mathbf{A}$ is a general matrix, one approach to finding $\mathbf{x}$ is to convert the problem to an equivalent one with an upper triangular coefficient matrix by systematically subtracting multiples of upper rows of the augmented matrix from lower rows (*elementary row operations*) so that more of the lower triangle elements become zero. This approach is called *row reduction* or *Gaussian elimination*. Because there are almost $n^2/2$ elements to convert to zero and each elementary row operation requires $\mathcal{O}(n)$ arithmetic operations, the whole process involves $\mathcal{O}(n^3)$ arithmetic operations.

Gauss elimination example

$$\text{With} \left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ -1 & 1 & -2 & -4 \\ 1 & 2 & -4 & -5 \end{array}\right),$$

we can subtract $-1/2$ times row 1 from row 2:

$$\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ 0 & 1 & -1/2 & 9/2 \\ 1 & 2 & -4 & -5 \end{array}\right)$$

and subtract $1/2$ times row 1 from row 3:

$$\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ 0 & 1 & -1/2 & 9/2 \\ 0 & 2 & -11/2 & -27/2 \end{array}\right)$$

then subtract 2 times row 2 from row 3:

$$\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ 0 & 1 & -1/2 & 9/2 \\ 0 & 0 & -9/2 & -45/2 \end{array}\right)$$

.

This gives an upper triangular coefficient matrix, and we can now find **x** by back substitution.

Practice problem

Use row reduction to write the following as an upper triangular system:

$$\left(\begin{array}{ccc|c} 3 & 1 & 5 & -60 \\ 4 & 1 & 0 & 37 \\ 3 & 4 & 4 & -76 \end{array}\right)$$

Row reduction algorithm

Inputs: $n \times n$ matrix $\mathbf{A}$ and $n \times 1$ $\mathbf{b}$.
Go from column 1 to column $n-1$ (denote the column index as c)
  Go from row $c+1$ to row n (denote the row index as r)
    Find the multiplier as $m_{r,c} = A_{r,c}/A_{c,c}$
    Apply a row operation: $A_{r,i} \leftarrow A_{r,i} - m_{r,c}A_{c,i}$ for $i = 1, 2, \ldots n$,
and $b_r \leftarrow b_r - m_{r,c}b_c$.

The result will be a new $\mathbf{A}$ (upper triangular) and $\mathbf{b}$, as long as none of the *pivots* $A_{c,c}$ are zero.

A modification: row pivoting

Row pivoting (also called *partial pivoting*) aims to avoid zero pivots, and also pivots that are almost zero (possibly due to roundoff; dividing by them could make the problem conditioning worse).
It relies on changing the order of rows, so that the pivot element is as far from zero as possible

Row reduction algorithm with pivoting

Inputs: $n \times n$ matrix $\mathbf{A}$ and $n \times 1$ $\mathbf{b}$.
Optional: start with a vector $\mathbf{v}$ of numbers 1 to n, or a permutation matrix $\mathbf{P}$ equal to $\mathbf{I}_n$
Go from column 1 to column $n-1$ (denote the column index as c)
  Let p be a row, out of the rows c to n, where $|A_{p,c}|$ is as large as possible
  Interchange rows p and c of $\mathbf{A}$, same for $\mathbf{b}$
  Optional: Also interchange rows p and c of $\mathbf{v}$ or $\mathbf{P}$
  Go from row $c+1$ to row n (denote the row index as r)
    Find the multiplier as $m_{r,c} = A_{r,c}/A_{c,c}$
    Apply a row operation: $A_{r,i} \leftarrow A_{r,i} - m_{r,c}A_{c,i}$ for $i = 1, 2, \ldots n$, and $b_r \leftarrow b_r - m_{r,c}b_c$.

Note: with this form of pivoting, all the multipliers m will be 1 or less in absolute value

Row pivoting example

With $\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ -1 & 1 & -2 & -4 \\ 1 & 2 & -4 & -5 \end{array}\right)$, no interchange is needed to begin with, since row 1 already has the largest absolute value of all rows for column 1. Subtract $-1/2$ times row 1 from row 2 and subtract $1/2$ times row 1 from row 3: $\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ 0 & 1 & -1/2 & 9/2 \\ 0 & 2 & -11/2 & -27/2 \end{array}\right)$. Now row 3 has the largest absolute value for column 2 across rows 2 and below, so interchange rows 2 and 3:
$\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ 0 & 2 & -11/2 & -27/2 \\ 0 & 1 & -1/2 & 9/2 \end{array}\right)$. Then subtract $1/2$ times row 2 from row 3: $\left(\begin{array}{ccc|c} 2 & 0 & 3 & 17 \\ 0 & 2 & -11/2 & -27/2 \\ 0 & 0 & 9/4 & 45/4 \end{array}\right)$.
This gives an upper triangular coefficient matrix, and we can now find $\mathbf{x}$ by back substitution.

Practice problem

Use row reduction with row pivoting to write the following as an upper triangular system:

$$\left(\begin{array}{ccc|c} 3 & 1 & 5 & -60 \\ 4 & 1 & 0 & 37 \\ 3 & 4 & 4 & -76 \end{array}\right)$$

Matrix factorization interpretation of row reduction

The effect of $\mathbf{A}$ of each elementary row operation is equivalent to multiplying it on the left by a unit lower triangular matrix whose only nonzero element in the lower triangle is $-m_{c,r}$ at position (c, r).
It turns out that when the original $\mathbf{A}$ is transformed to an upper triangular matrix $\mathbf{U}$, we can recover $\mathbf{A}$ by multiplying $\mathbf{U}$ by a unit lower triangular matrix $\mathbf{L}$, each of whose lower triangular elements is equal to the corresponding $m_{c,r}$ multiplier.
So $\mathbf{A} = \mathbf{LU}$, where the matrix factors $\mathbf{L}, \mathbf{U}$ can be found using row reduction. This is called the *LU factorization/decomposition* of $\mathbf{A}$.

Example

For $\mathbf{A} = \begin{pmatrix} 2 & 0 & 3 \\ -1 & 1 & -2 \\ 1 & 2 & -4 \end{pmatrix}$,

$$\mathbf{U} = \begin{pmatrix} 2 & 0 & 3 \\ 0 & 1 & -1/2 \\ 0 & 0 & -9/2 \end{pmatrix}, \mathbf{L} = \begin{pmatrix} 1 & 0 & 0 \\ -1/2 & 1 & 0 \\ 1/2 & 2 & 1 \end{pmatrix}.$$

We can verify that $\mathbf{U}$ is upper triangular, $\mathbf{L}$ is lower triangular, and $\mathbf{LU} = \mathbf{A}$.

Factorization using row reduction with pivoting

To avoid zero (or small) pivots, we usually want to add row pivoting to the row reduction. The result is then $\mathbf{L}$ and $\mathbf{U}$ factors that multiply to $\mathbf{A}$ with the same rows interchanged as done in the pivoting steps, that is $\mathbf{LU} = \mathbf{PA}$.

Example with row pivoting

For $\mathbf{A} = \begin{pmatrix} 2 & 0 & 3 \\ -1 & 1 & -2 \\ 1 & 2 & -4 \end{pmatrix}$,

$$\mathbf{U} = \begin{pmatrix} 2 & 0 & 3 \\ 0 & 2 & -11/2 \\ 0 & 0 & 9/4 \end{pmatrix}, \mathbf{L} = \begin{pmatrix} 1 & 0 & 0 \\ -1/2 & 1 & 0 \\ 1/2 & 1/2 & 1 \end{pmatrix}.$$

We can verify that $\mathbf{U}$ is upper triangular, $\mathbf{L}$ is lower triangular, and

$$\mathbf{LU} = \mathbf{PA} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 2 & 0 & 3 \\ -1 & 1 & -2 \\ 1 & 2 & -4 \end{pmatrix} = \begin{pmatrix} 2 & 0 & 3 \\ 1 & 2 & -4 \\ -1 & 1 & -2 \end{pmatrix}.$$

Uniqueness of solution to a linear system

If Gauss elimination with row pivoting returns a solution, it normally will be a unique solution, corresponding to $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$. If $\mathbf{A}$ has no inverse (is *singular*), there is no unique solution and the Gauss elimination + back substitution will fail (there will be zero or infinitely many solutions).

Accuracy of solution where there are errors

If **A** or **b** have small fractional errors, or if there is roundoff error in the row reduction computations, the computed **x** may differ from the true one. The expected fractional error is proportional to the *condition number* of **A** times the roundoff machine epsilon (or magnitude of other fractional error) ϵ. The condition number is at least 1, goes to infinity where **A** has no inverse, and becomes very large if **A** does have an inverse but is close to a singular matrix.

Solving linear systems with LU factorization

If $\mathbf{PA} = \mathbf{LU}$, we can solve any linear system $\mathbf{Ax} = \mathbf{b}$ in two steps:
1) Solve $\mathbf{Ly} = \mathbf{Pb}$ for $\mathbf{y}$ (forward substitution)
2) Solve $\mathbf{Ux} = \mathbf{y}$ for $\mathbf{x}$ (back substitution).
Each step only requires n^2 operations, as opposed to $\mathcal{O}(n^3)$ to get the $\mathbf{L}$ and $\mathbf{U}$ factors using row reduction.

Finding the matrix determinant with LU factorization

- The determinant of a matrix product is the product of the determinants of the factors.
- The determinant of a triangular matrix is the product of the elements on the main diagonal.
- Each interchange of rows in a matrix flips the determinant's sign.

Finding the matrix inverse with LU factorization

- The inverse of $\mathbf{A}$ can be found by solving $\mathbf{AX} = \mathbf{I}$ for $\mathbf{X}$.
- In theory, for any invertible $\mathbf{A}$, we can write the solution to $\mathbf{Ax} = \mathbf{b}$ as $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$. However, finding the inverse is <u>not</u> the most computationally accurate or efficient way to solve the problem – it's better to use the LU factorization of $\mathbf{A}$ to set up and solve triangular systems.