MA4230 Matrix Computation

Chapter 4: Linear Systems and Least Squares Problems

4.1 Gaussian elimination: LU factorization

The problem

Given $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$, find $x \in \mathbb{R}^n$ s.t.

$$Ax = b.$$

$\Longrightarrow$ Solve by Gaussian elimination.

No/partial/full pivoting $\Longrightarrow A = LU \;/\; PA = LU \;/\; PAQ = LU$

We start with the LU factorization (Gauß without pivoting).

Definition (Lower-triangular and unit lower-triangular matrices)

A matrix $L \in \mathbb{R}^{n \times n}$ is called **lower-triangular** iff L^{T} is upper-triangular.
A matrix $L \in \mathbb{R}^{n \times n}$ is called **unit lower-triangular** iff L is lower-triangular and all of its diagonal entries are equal to 1.

LU factorization

Definition (LU factorization)

Let $A \in \mathbb{R}^{n \times n}$. If $\exists\ L \in \mathbb{R}^{n \times n}$ lower-triangular, $U \in \mathbb{R}^{n \times n}$ upper-triangular s.t. $A = LU$, then this factorization is called a **LU factorization** of A.

Gaussian elimination transforms A into an upper-triangular matrix

$$U = L_{n-1} \cdots L_2 L_1 A \in \mathbb{R}^{n \times n}$$

with $L_1, \ldots, L_{n-1} \in \mathbb{R}^{n \times n}$ unit lower-triangular and of the form

$$L_1 = \begin{pmatrix} 1 & & & \\ * & 1 & & \\ \vdots & & \ddots & \\ \vdots & & & \ddots & \\ * & & & & 1 \end{pmatrix}, L_2 = \begin{pmatrix} 1 & & & & \\ & 1 & & & \\ & * & \ddots & & \\ & \vdots & & \ddots & \\ & * & & & 1 \end{pmatrix}, \cdots, L_{n-1} = \begin{pmatrix} 1 & & & & \\ & 1 & & & \\ & & \ddots & & \\ & & & 1 & \\ & & & * & 1 \end{pmatrix}$$

with zero-entries not shown.

Assuming this is possible, obtain $A = LU$ with $L := L_1^{-1} \cdots L_{n-1}^{-1} \in \mathbb{R}^{n \times n}$ unit lower-triangular (exercise) and $U \in \mathbb{R}^{n \times n}$ upper-triangular.

Gaussian elimination: Example

Consider $A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 1 & 1 & 2 & -2 \\ -1 & 4 & -1 & 1 \\ 1 & 3 & -3 & 4 \end{pmatrix}$. We illustrate Gaussian elimination.

L_1: The first step is to eliminate the sub-diagonal entries in the first column of A via adding $\frac{1}{2}/-\frac{1}{2}/\frac{1}{2}$ times row 1 to row 2/3/4:

$$L_1 A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 2 & 5/2 & -5/2 \\ 0 & 3 & -3/2 & 3/2 \\ 0 & 4 & -5/2 & 7/2 \end{pmatrix} \quad \text{with} \quad L_1 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1/2 & 1 & 0 & 0 \\ -1/2 & 0 & 1 & 0 \\ 1/2 & 0 & 0 & 1 \end{pmatrix}.$$

L_2: The second step is to eliminate the sub-diagonal entries in the second column of $L_1 A$ via adding $-\frac{3}{2}/-2$ times row 2 to row 3/4:

$$L_2 L_1 A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 2 & 5/2 & -5/2 \\ 0 & 0 & -21/4 & 21/4 \\ 0 & 0 & -15/2 & 17/2 \end{pmatrix} \quad \text{with} \quad L_2 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -\frac{3}{2} & 1 & 0 \\ 0 & -2 & 0 & 1 \end{pmatrix}.$$

L_2: The second step is to eliminate the sub-diagonal entries in the second column of $L_1 A$ via adding $-\frac{3}{2}/-2$ times row 2 to row 3/4:

$$L_2 L_1 A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 2 & 5/2 & -5/2 \\ 0 & 0 & -21/4 & 21/4 \\ 0 & 0 & -15/2 & 17/2 \end{pmatrix} \quad \text{with} \quad L_2 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -\frac{3}{2} & 1 & 0 \\ 0 & -2 & 0 & 1 \end{pmatrix}.$$

L_3: The third step is to eliminate the sub-diagonal entries in the third column of $L_2 L_1 A$ via adding $-\frac{10}{7}$ times row 3 to row 4:

$$L_3 L_2 L_1 A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 2 & 5/2 & -5/2 \\ 0 & 0 & -21/4 & 21/4 \\ 0 & 0 & 0 & 1 \end{pmatrix} =: U \text{ with } L_3 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -\frac{10}{7} & 1 \end{pmatrix}.$$

We find that $A = LU$ with U as above and L given by

$$L := L_1^{-1} L_2^{-1} L_3^{-1}$$

is a LU factorization of A. Indeed, let's compute L:

$$L = L_1^{-1} L_2^{-1} L_3^{-1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ \frac{1}{2} & 1 & 0 & 0 \\ -\frac{1}{2} & 0 & 1 & 0 \\ \frac{1}{2} & 0 & 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -\frac{3}{2} & 1 & 0 \\ 0 & -2 & 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -\frac{10}{7} & 1 \end{pmatrix}^{-1}$$

$$= \begin{pmatrix} 1 & 0 & 0 & 0 \\ -\frac{1}{2} & 1 & 0 & 0 \\ \frac{1}{2} & 0 & 1 & 0 \\ -\frac{1}{2} & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & \frac{3}{2} & 1 & 0 \\ 0 & 2 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{10}{7} & 1 \end{pmatrix}$$

$$= \begin{pmatrix} 1 & 0 & 0 & 0 \\ -\frac{1}{2} & 1 & 0 & 0 \\ \frac{1}{2} & \frac{3}{2} & 1 & 0 \\ -\frac{1}{2} & 2 & \frac{10}{7} & 1 \end{pmatrix}$$

Note how simple it is to compute L: the matrices L_i can be inverted by negating their sub-diagonal entries, and the matrix L can be obtained by collecting these values appropriately. Coincidence? No:

Generally, if the i-th column x_i of the matrix $L_{i-1}\cdots L_1 A$ (the matrix A if $i=1$) is the vector $x_i = (x_{1i},\dots,x_{ni})^{\mathrm{T}}$, then

$$L_i = \begin{pmatrix} 1 & & & & \\ & \ddots & & & \\ & & 1 & & \\ & & -\frac{x_{i+1,i}}{x_{ii}} & 1 & \\ & & \vdots & & \ddots & \\ & & -\frac{x_{ni}}{x_{ii}} & & & 1 \end{pmatrix} = I_n - l_i e_i^{\mathrm{T}} \in \mathbb{R}^{n\times n}, \quad l_i := \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \frac{x_{i+1,i}}{x_{ii}} \\ \vdots \\ \frac{x_{ni}}{x_{ii}} \end{pmatrix} \in \mathbb{R}^n.$$

- $L_i^{-1} = I_n + l_i e_i^{\mathrm{T}}$: $(I_n - l_i e_i^{\mathrm{T}})(I_n + l_i e_i^{\mathrm{T}}) = I_n - l_i e_i^{\mathrm{T}} l_i e_i^{\mathrm{T}} = I_n - \langle l_i, e_i\rangle l_i e_i^{\mathrm{T}} = I_n$.
- $L = L_1^{-1}\cdots L_{n-1}^{-1}$ is given by

$$L = \begin{pmatrix} 1 & & & & \\ \frac{x_{21}}{x_{11}} & 1 & & & \\ \frac{x_{31}}{x_{11}} & \frac{x_{32}}{x_{22}} & 1 & & \\ \vdots & \vdots & \ddots & \ddots & \\ \frac{x_{n1}}{x_{11}} & \frac{x_{n2}}{x_{22}} & \cdots & \frac{x_{n,n-1}}{x_{n-1,n-1}} & 1 \end{pmatrix}.$$

Indeed, looking at the product of two such matrices we find

$$L_i^{-1} L_{i+1}^{-1} = (I_n + l_i e_i^{\mathrm{T}})(I_n + l_{i+1} e_{i+1}^{\mathrm{T}}) = I_n + l_i e_i^{\mathrm{T}} + l_{i+1} e_{i+1}^{\mathrm{T}}.$$

Gaussian elimination without pivoting: Algorithm

Given $A \in \mathbb{R}^{n\times n}$, do as follows:

$L = I_n$, $U = A$
for $i = 1, \ldots, n-1$ **do**
  for $j = i+1, \ldots, n$ **do**
    $l_{ji} = \frac{u_{ji}}{u_{ii}}$
    $u_{j,i:n} = u_{j,i:n} - l_{ji} u_{i,i:n}$
  end for
end for.

Warning: A needs to be such that no division by zero happens.

Theorem

The above algorithm requires $\sim \frac{2}{3} n^3$ *flops.*

Proof: Exercise.

Compare with $\sim \frac{4}{3} n^3$ flops for QR via Householder.

Solving linear systems via LU

Problem: Given $A \in \mathbb{R}^{n\times n}$, $b \in \mathbb{R}^n$, find $x \in \mathbb{R}^n$ s.t. $Ax = b$.

If there exists a LU factorization $A = LU$, we have

$$Ax = b \iff LUx = b \iff \begin{cases} Ly = b, \\ Ux = y. \end{cases}$$

Therefore, once a LU factorization is computed ($\mathcal{O}(n^3)$ flops), we can first solve $Ly = b$ for y by forward substitution ($\mathcal{O}(n^2)$ flops) and then $Ux = y$ for x by backward substitution ($\mathcal{O}(n^2)$ flops).

$\Longrightarrow$ But does every matrix have a LU factorization? Unfortunately, no.

E.g., $A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$ does not have a LU factorization. Indeed, if there were $L = \begin{pmatrix} l_{11} & 0 \\ l_{21} & l_{22} \end{pmatrix} \in \mathbb{R}^{2\times 2}$ and $U = \begin{pmatrix} u_{11} & u_{12} \\ 0 & u_{22} \end{pmatrix} \in \mathbb{R}^{2\times 2}$ such that $A = LU$, then $l_{11}u_{11} = 0$ and $l_{11}u_{12} = l_{21}u_{11} = l_{21}u_{12} + l_{22}u_{22} = 1$, contradiction.

Gaussian elimination in its current form (without pivoting) is impractical to solve general linear systems. For instance, it fails for the matrix

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$$

due to division by zero in the first step.

More dramatically, the algorithm is not stable for general $n \times n$ matrices as we will see later in this course.

Improvement in stability via pivoting $\Longrightarrow$ next section.

4.2 Gaussian elimination with partial pivoting: PA=LU factorization

How to improve Gaussian elimination? Key observation

In i-th step of Gauß, add multiples of row i to rows $i+1, \dots, n$ to obtain

$$\left(\begin{array}{ccccccc} x_{11} & x_{12} & \cdots & x_{1i} & x_{1,i+1} & \cdots & x_{1n} \\ & x_{22} & \cdots & x_{2i} & x_{2,i+1} & \cdots & x_{2n} \\ & & \ddots & \vdots & \vdots & & \vdots \\ \hline & & & x_{ii} & x_{i,i+1} & \cdots & x_{in} \\ & & & x_{i+1,i} & x_{i+1,i+1} & \cdots & x_{i+1,n} \\ & & & \vdots & \vdots & & \vdots \\ & & & x_{ni} & x_{n,i+1} & \cdots & x_{nn} \end{array}\right) \Longrightarrow \left(\begin{array}{ccccccc} x_{11} & x_{12} & \cdots & x_{1i} & x_{1,i+1} & \cdots & x_{1n} \\ & x_{22} & \cdots & x_{2i} & x_{2,i+1} & \cdots & x_{2n} \\ & & \ddots & \vdots & \vdots & & \vdots \\ \hline & & & x_{ii} & x_{i,i+1} & \cdots & x_{in} \\ & & & 0 & * & \cdots & * \\ & & & \vdots & \vdots & & \vdots \\ & & & 0 & * & \cdots & * \end{array}\right)$$

We call $x_{ii} \neq 0$ the **pivot**. Observation: Instead, can also add multiples of row j with some $j \in \{i+1, \dots, n\}$ such that $x_{ji} \neq 0$ to rows $i, \dots, j-1, j+1, \dots, n$ to create zeros as follows:

$$\left(\begin{array}{ccccccc} x_{11} & x_{12} & \cdots & x_{1i} & x_{1,i+1} & \cdots & x_{1n} \\ & x_{22} & \cdots & x_{2i} & x_{2,i+1} & \cdots & x_{2n} \\ & & \ddots & \vdots & \vdots & & \vdots \\ \hline & & & x_{ii} & x_{i,i+1} & \cdots & x_{in} \\ & & & \vdots & \vdots & & \vdots \\ & & & x_{ji} & x_{j,i+1} & \cdots & x_{jn} \\ & & & \vdots & \vdots & & \vdots \\ & & & x_{ni} & x_{n,i+1} & \cdots & x_{nn} \end{array}\right) \Longrightarrow \left(\begin{array}{ccccccc} x_{11} & x_{12} & \cdots & x_{1i} & x_{1,i+1} & \cdots & x_{1n} \\ & x_{22} & \cdots & x_{2i} & x_{2,i+1} & \cdots & x_{2n} \\ & & \ddots & \vdots & \vdots & & \vdots \\ \hline & & & 0 & * & \cdots & * \\ & & & \vdots & \vdots & & \vdots \\ & & & 0 & * & \cdots & * \\ & & & x_{ji} & x_{j,i+1} & \cdots & x_{jn} \\ & & & 0 & * & \cdots & * \\ & & & \vdots & \vdots & & \vdots \\ & & & 0 & * & \cdots & * \end{array}\right).$$

In this case, $x_{ji} \neq 0$ is called the pivot.

Gaussian elimination with partial pivoting

This procedure is thought of as follows:

In the i-th step,

1. choose a pivot $x_{ji} \neq 0$ from column i and row j (some $j \in \{i, \ldots, n\}$),
2. permute the rows such that x_{ji} is moved to the main diagonal,
3. do a standard Gaussian elimination step.

For numerical stability, the pivot is chosen as the largest entry in modulus in column i and rows $i, \ldots, n$.

This is called **Gaussian elimination with partial pivoting** and leads to a LU factorization of PA for some permutation matrix P.

PA=LU factorization

Definition (PA=LU factorization)

Let $A \in \mathbb{R}^{n\times n}$. If there exist a lower-triangular matrix $L \in \mathbb{R}^{n\times n}$, an upper-triangular matrix $U \in \mathbb{R}^{n\times n}$, and a permutation matrix $P \in \mathbb{R}^{n\times n}$ (i.e., a matrix which has exactly one entry 1 in each row and column and zeros elsewhere) s.t.

$$PA = LU,$$

then we call this factorization a **PA=LU factorization** or a **LU factorization with partial pivoting** corresponding to A.

Remark: Permutation matrices are orthogonal matrices.

Gaussian elimination with partial pivoting: Example

Consider $A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 1 & 1 & 2 & -2 \\ -1 & 4 & -1 & 1 \\ 1 & 3 & -3 & 4 \end{pmatrix}$.

P_1: As $\max\{|-2|, |1|, |-1|, |1|\} = |-2|$, choose the $(1,1)$-entry as pivot. Since this is already on the diagonal, no permutation is needed:

$$P_1 A = A \quad \text{with} \quad P_1 := I_4.$$

L_1: Eliminate sub-diagonal entries in first column of $P_1 A = A$ via adding $\frac{1}{2}/-\frac{1}{2}/\frac{1}{2}$ times row 1 to row 2/3/4:

$$L_1 P_1 A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 2 & 5/2 & -5/2 \\ 0 & 3 & -3/2 & 3/2 \\ 0 & 4 & -5/2 & 7/2 \end{pmatrix} \quad \text{with} \quad L_1 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1/2 & 1 & 0 & 0 \\ -1/2 & 0 & 1 & 0 \\ 1/2 & 0 & 0 & 1 \end{pmatrix}.$$

Recall from previous slide: $L_1P_1A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 2 & 5/2 & -5/2 \\ 0 & 3 & -3/2 & 3/2 \\ 0 & 4 & -5/2 & 7/2 \end{pmatrix}$.

P_2: As $\max\{|2|, |3|, |4|\} = |4|$, choose the $(4,2)$-entry as pivot. To this end, we permute rows 2 and 4:

$$P_2L_1P_1A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 4 & -5/2 & 7/2 \\ 0 & 3 & -3/2 & 3/2 \\ 0 & 2 & 5/2 & -5/2 \end{pmatrix} \text{ with } P_2 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

L_2: Eliminate sub-diagonal entries in second column of $P_2L_1P_1A$ via adding $-\frac{3}{4}/-\frac{1}{2}$ times row 2 to row 3/4:

$$L_2P_2L_1P_1A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 4 & -5/2 & 7/2 \\ 0 & 0 & 3/8 & -9/8 \\ 0 & 0 & 15/4 & -17/4 \end{pmatrix} \text{ with } L_2 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -3/4 & 1 & 0 \\ 0 & -1/2 & 0 & 1 \end{pmatrix}$$

Recall from previous slide: $L_2P_2L_1P_1A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 4 & -5/2 & 7/2 \\ 0 & 0 & 3/8 & -9/8 \\ 0 & 0 & 15/4 & -17/4 \end{pmatrix}$.

P_3: As $\max\{|\frac{3}{8}|, |\frac{15}{4}|\} = |\frac{15}{4}|$, choose the $(4,3)$-entry as pivot. To this end, we permute rows 3 and 4:

$$P_3L_2P_2L_1P_1A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 4 & -5/2 & 7/2 \\ 0 & 0 & 15/4 & -17/4 \\ 0 & 0 & 3/8 & -9/8 \end{pmatrix} \text{ with } P_3 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

L_3: Eliminate sub-diagonal entries in third column of $P_3L_2P_2L_1P_1A$ via adding $-\frac{1}{10}$ times row 3 to row 4:

$$L_3P_3L_2P_2L_1P_1A = \begin{pmatrix} -2 & 2 & 1 & -1 \\ 0 & 4 & -5/2 & 7/2 \\ 0 & 0 & 15/4 & -17/4 \\ 0 & 0 & 0 & -7/10 \end{pmatrix} =: U, \; L_3 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -\frac{1}{10} & 1 \end{pmatrix}$$

$\implies \quad L_3P_3L_2P_2L_1P_1A = U$. How to obtain from this a PA=LU factn?

Set $L_3' := L_3$, $L_2' := P_3L_2P_3^{-1}$, and $L_1' := P_3P_2L_1P_2^{-1}P_3^{-1}$, i.e.,

$$L_3' = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -\frac{1}{10} & 1 \end{pmatrix}, L_2' = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -1/2 & 1 & 0 \\ 0 & -3/4 & 0 & 1 \end{pmatrix}, L_1' = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1/2 & 1 & 0 & 0 \\ 1/2 & 0 & 1 & 0 \\ -1/2 & 0 & 0 & 1 \end{pmatrix}$$

Then, $L_3'L_2'L_1'P_3P_2P_1A = L_3P_3L_2P_2L_1P_1A = U$.

We find that $PA = LU$ with

$$P := P_3P_2P_1 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \; L := (L_3'L_2'L_1')^{-1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -1/2 & 1 & 0 & 0 \\ -1/2 & 1/2 & 1 & 0 \\ 1/2 & 3/4 & 1/10 & 1 \end{pmatrix}$$

is a PA=LU factorization.

More generally, . . .

Gaussian elimination with partial pivoting transforms $A \in \mathbb{R}^{n\times n}$ into an upper-triangular $U \in \mathbb{R}^{n\times n}$ by Gaussian elimination with an additional left-multiplication of a permutation matrix P_i at the beginning of step i:

$$L_{n-1}P_{n-1}\cdots L_2P_2L_1P_1A = U.$$

Here, $P_1, \ldots, P_{n-1} \in \mathbb{R}^{n\times n}$ are permutation matrices and $L_1, \ldots, L_{n-1} \in \mathbb{R}^{n\times n}$ are unit lower-triangular.

Set $L'_{n-1} := L_{n-1}$, $L'_i := P_{n-1}\cdots P_{i+1}L_iP_{i+1}^{-1}\cdots P_{n-1}^{-1}$ for $1 \le i \le n-2$.

$$\implies \qquad (L'_{n-1}\cdots L'_2L'_1)(P_{n-1}\cdots P_2P_1)A = U.$$

Observe that the matrix L'_i has the same structure as L_i. We then obtain that $PA = LU$ is a PA=LU factorization corresponding to A with

$$L := (L'_{n-1}\cdots L'_2L'_1)^{-1}, \qquad P := P_{n-1}\cdots P_2P_1.$$

Note P is a permutation matrix as product of permutation matrices, and that L is well-defined and lower-triangular.

Gaussian elimination with partial pivoting: Algorithm

Given $A \in \mathbb{R}^{n \times n}$, do as follows:

$P = I_n$, $L = I_n$, $U = A$
for $i = 1, \ldots, n-1$ **do**
  Choose $r \in \{i, \ldots, n\}$ such that $|u_{ri}| = \max_{k \in \{i,\ldots,n\}} |u_{ki}|$
  $u_{i,i:n} \leftrightarrow u_{r,i:n}$
  $l_{i,1:i-1} \leftrightarrow l_{r,1:i-1}$
  $p_{i,1:n} \leftrightarrow p_{r,1:n}$
  for $j = i+1, \ldots, n$ **do**
    $l_{ji} = \frac{u_{ji}}{u_{ii}}$
    $u_{j,i:n} = u_{j,i:n} - l_{ji} u_{i,i:n}$
  end for
end for.

Here, "$\leftrightarrow$" denotes "interchange".

Warning: A needs to be such that no division by zero happens in the algorithm above (as an exercise, think about how to obtain a PA=LU factorization if all candidates for pivots are zero at some step i).

Work of Gauß with partial pivoting

- pivot selection requires $\mathcal{O}(n^2)$ operations overall.

$\Longrightarrow$ To leading order, Gauß with partial pivoting requires same amount of flops as Gauß without pivoting, i.e., $\sim \frac{2}{3}n^3$.

Gaussian elimination with partial pivoting is the standard way to solve linear systems on a computer.

Solving linear systems via PA=LU factorization

Problem: Given $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$, find $x \in \mathbb{R}^n$ s.t. $Ax = b$.

If there exists a PA=LU factorization $PA = LU$, we have

$$Ax = b \iff PAx = Pb \iff LUx = Pb \iff \begin{cases} Ly = Pb, \\ Ux = y. \end{cases}$$

Therefore, once a PA=LU factorization is computed ($\mathcal{O}(n^3)$ flops), we can first form $\tilde{b} := Pb$, then solve $Ly = \tilde{b}$ by forward substitution ($\mathcal{O}(n^2)$ flops) and then $Ux = y$ for x by backward substitution ($\mathcal{O}(n^2)$ flops).

What about existence of LU and PA=LU factorization?
(Recall we already know that not every matrix has a LU factorization.)

Theorem (Existence of LU and PA=LU factorization)

(i) *Any matrix $A \in \mathbb{R}^{n \times n}$ has a PA=LU factorization.*

(ii) *Let $A \in \mathbb{R}^{n \times n}$ invertible. Then, there exists a LU factorization of A iff $\det(A_{1:i,1:i}) \neq 0$ for all $1 \leq i \leq n$.*

4.3 Gaussian elimination with full pivoting: PAQ=LU factorization

Full pivoting: A further improvement in stability

Idea: Every entry of the sub-matrix $X_{i:n,i:n}$ of the working matrix X at step i is a candidate for the pivot.

Rk: *Full pivoting is rarely used in practice due to large computational cost.*

Gaussian elimination with full pivoting leads to a PAQ=LU factorization:

Definition (PAQ=LU factorization)

Let $A \in \mathbb{R}^{n \times n}$. If there exist a lower-triangular matrix $L \in \mathbb{R}^{n \times n}$, an upper-triangular matrix $U \in \mathbb{R}^{n \times n}$, and permutation matrices $P, Q \in \mathbb{R}^{n \times n}$ such that there holds

$$PAQ = LU,$$

then we call this a **PAQ=LU factorization** or a **LU factorization with full pivoting** corresponding to A.

Note: Any matrix $A \in \mathbb{R}^{n \times n}$ admits a PAQ=LU factorization with $Q = I_n$.

Example: Gaussian elimination with full pivoting

Consider $A := \begin{pmatrix} -2 & 2 & 1 & -1 \\ 1 & 1 & 2 & -2 \\ -1 & 4 & -1 & 1 \\ 1 & 3 & -3 & 4 \end{pmatrix}$.

P_1, Q_1: As $\max_{i,j\in\{1,\dots,4\}} |a_{ij}| = |4|$, we choose the $(3,2)$-entry 4 as pivot (note we could have also chosen the $(4,4)$-entry 4).
To this end, we permute columns 1 and 2, and then rows 1 and 3:

$$P_1 A Q_1 = \begin{pmatrix} 4 & -1 & -1 & 1 \\ 1 & 1 & 2 & -2 \\ 2 & -2 & 1 & -1 \\ 3 & 1 & -3 & 4 \end{pmatrix}, \quad Q_1 := \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, P_1 := \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

L_1: We eliminate the sub-diagonal entries in the first column of $P_1 A Q_1$ via adding $-\frac{1}{4}/-\frac{1}{2}/-\frac{3}{4}$ times row 1 to row 2/3/4:

$$L_1 P_1 A Q_1 = \begin{pmatrix} 4 & -1 & -1 & 1 \\ 0 & 5/4 & 9/4 & -9/4 \\ 0 & -3/2 & 3/2 & -3/2 \\ 0 & 7/4 & -9/4 & 13/4 \end{pmatrix} \quad \text{with } L_1 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ -1/4 & 1 & 0 & 0 \\ -1/2 & 0 & 1 & 0 \\ -3/4 & 0 & 0 & 1 \end{pmatrix}$$

$$L_1P_1AQ_1 = \begin{pmatrix} 4 & -1 & -1 & 1 \\ 0 & 5/4 & 9/4 & -9/4 \\ 0 & -3/2 & 3/2 & -3/2 \\ 0 & 7/4 & -9/4 & 13/4 \end{pmatrix} =: A_1.$$

P_2, Q_2: As $\max\{|\frac{5}{4}|, |-\frac{3}{2}|, |\frac{7}{4}|, |\frac{9}{4}|, |\frac{3}{2}|, |-\frac{9}{4}|, |-\frac{9}{4}|, |-\frac{3}{2}|, |\frac{13}{4}|\} = |\frac{13}{4}|$, we choose the $(4,4)$-entry $\frac{13}{4}$ as pivot.
To this end, we permute columns 2 and 4, and then rows 2 and 4:

$$P_2A_1Q_2 = \begin{pmatrix} 4 & 1 & -1 & -1 \\ 0 & 13/4 & -9/4 & 7/4 \\ 0 & -3/2 & 3/2 & -3/2 \\ 0 & -9/4 & 9/4 & 5/4 \end{pmatrix}, \quad Q_2 := P_2 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

L_2: We eliminate the sub-diagonal entries in the second column of $P_2A_1Q_2$ via adding $\frac{6}{13}/\frac{9}{13}$ times row 2 to row 3/4:

$$L_2P_2A_1Q_2 = \begin{pmatrix} 4 & 1 & -1 & -1 \\ 0 & 13/4 & -9/4 & 7/4 \\ 0 & 0 & 6/13 & -9/13 \\ 0 & 0 & 9/13 & 32/13 \end{pmatrix} \quad \text{with } L_2 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 6/13 & 1 & 0 \\ 0 & 9/13 & 0 & 1 \end{pmatrix}$$

$$L_2P_2A_1Q_2 = \begin{pmatrix} 4 & 1 & -1 & -1 \\ 0 & 13/4 & -9/4 & 7/4 \\ 0 & 0 & 6/13 & -9/13 \\ 0 & 0 & 9/13 & 32/13 \end{pmatrix} =: A_2.$$

P_3, Q_3: As $\max\{|\frac{6}{13}|, |\frac{9}{13}|, |-\frac{9}{13}|, |\frac{32}{13}|\} = |\frac{32}{13}|$, we choose the $(4,4)$-entry $\frac{32}{13}$ as pivot.

To this end, we permute columns 3 and 4, and then rows 3 and 4:

$$P_3A_2Q_3 = \begin{pmatrix} 4 & 1 & -1 & -1 \\ 0 & 13/4 & 7/4 & -9/4 \\ 0 & 0 & 32/13 & 9/13 \\ 0 & 0 & -9/13 & 6/13 \end{pmatrix}, \quad Q_3 := P_3 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

L_3: We eliminate the sub-diagonal entries in the third column of $P_3A_2Q_3$ via adding $\frac{9}{32}$ times row 3 to row 4:

$$L_3P_3A_2Q_3 = \begin{pmatrix} 4 & 1 & -1 & -1 \\ 0 & 13/4 & 7/4 & -9/4 \\ 0 & 0 & 32/13 & 9/13 \\ 0 & 0 & 0 & 21/32 \end{pmatrix} =: U \text{ with } L_3 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{9}{32} & 1 \end{pmatrix}$$

$\Longrightarrow$ $L_3P_3L_2P_2L_1P_1AQ_1Q_2Q_3 = U$ is upper-triangular.

Have $L_3'L_2'L_1'P_3P_2P_1AQ_1Q_2Q_3 = L_3P_3L_2P_2L_1P_1AQ_1Q_2Q_3 = U$ with $L_3' := L_3$,

$$L_2' := P_3L_2P_3^{-1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 9/13 & 1 & 0 \\ 0 & 6/13 & 0 & 1 \end{pmatrix}, \quad L_1' := P_3P_2L_1P_2^{-1}P_3^{-1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -3/4 & 1 & 0 & 0 \\ -1/4 & 0 & 1 & 0 \\ -1/2 & 0 & 0 & 1 \end{pmatrix}.$$

$\Longrightarrow$ We find that $PAQ = LU$ with

$$P := P_3P_2P_1 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \quad Q := Q_1Q_2Q_3 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix},$$

$$L := (L_3'L_2'L_1')^{-1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 3/4 & 1 & 0 & 0 \\ 1/4 & -9/13 & 1 & 0 \\ 1/2 & -6/13 & -9/32 & 1 \end{pmatrix}, \quad U = \begin{pmatrix} 4 & 1 & -1 & -1 \\ 0 & 13/4 & 7/4 & -9/4 \\ 0 & 0 & 32/13 & 9/13 \\ 0 & 0 & 0 & 21/32 \end{pmatrix}$$

is a PAQ = LU factorization.

More generally, ...

Gaussian elimination with full pivoting transforms $A \in \mathbb{R}^{n\times n}$ into an upper-triangular $U \in \mathbb{R}^{n\times n}$ by Gaussian elimination with an additional right-multiplication of a permutation matrix Q_i and left-multiplication of a permutation matrix P_i at the beginning of step i:

$$L_{n-1}P_{n-1}\cdots L_2P_2L_1P_1AQ_1Q_2\cdots Q_{n-1} = U.$$

Here, $P_1,\dots,P_{n-1},Q_1,\dots,Q_{n-1} \in \mathbb{R}^{n\times n}$ are permutation matrices and $L_1,\dots,L_{n-1} \in \mathbb{R}^{n\times n}$ are unit lower-triangular.

We deduce that

$$(L'_{n-1}\cdots L'_2L'_1)(P_{n-1}\cdots P_2P_1)A(Q_1Q_2\cdots Q_{n-1}) = U$$

with $L'_{n-1} := L_{n-1}$ and $L'_i := P_{n-1}\cdots P_{i+1}L_iP_{i+1}^{-1}\cdots P_{n-1}^{-1}$ for $i \in \{1,\dots,n-2\}$. We then obtain that $PAQ = LU$ is a PAQ=LU factorization corresponding to A with

$$L := (L'_{n-1}\cdots L'_2L'_1)^{-1}, \quad P := P_{n-1}\cdots P_2P_1, \quad Q := Q_1Q_2\cdots Q_{n-1}.$$

Note that P and Q are permutation matrices as products of permutation matrices, and that L is well-defined and lower-triangular.

Advantages and drawbacks

+ Full pivoting further improves stability compared to partial pivoting.
– Pivot selection for full pivoting requires $\mathcal{O}(n^3)$ operations overall.

As an exercise, think about how a PAQ=LU factorization can be used to solve a linear system $Ax = b$.

4.4 Symmetric Gaussian elimination: Cholesky factorization

Definite matrices

Definition (positive/negative (semi)definiteness)

A symmetric matrix $A \in \mathbb{R}^{n \times n}$ is called

(i) **positive definite**, denoted $A \succ 0$, iff

$$\langle x, Ax \rangle = x^{\mathrm{T}} Ax > 0 \quad \forall x \in \mathbb{R}^n \backslash \{0\}.$$

(ii) **positive semidefinite**, denoted $A \succeq 0$, iff

$$\langle x, Ax \rangle = x^{\mathrm{T}} Ax \geq 0 \quad \forall x \in \mathbb{R}^n.$$

(iii) **negative definite**, denoted $A \prec 0$, iff

$$\langle x, Ax \rangle = x^{\mathrm{T}} Ax < 0 \quad \forall x \in \mathbb{R}^n \backslash \{0\}.$$

(iv) **negative semidefinite**, denoted $A \preceq 0$, iff

$$\langle x, Ax \rangle = x^{\mathrm{T}} Ax \leq 0 \quad \forall x \in \mathbb{R}^n.$$

An equivalent characterization via eigenvalues

Theorem (Characterization of positive/negative (semi)definite matrices)

For a symmetric matrix $A \in \mathbb{R}^{n\times n}$, we have

(i) $A \succ 0 \iff$ *all eigenvalues of A are positive*,

(ii) $A \succeq 0 \iff$ *all eigenvalues of A are non-negative*,

(iii) $A \prec 0 \iff$ *all eigenvalues of A are negative*,

(iv) $A \preceq 0 \iff$ *all eigenvalues of A are non-positive*.

Proof: Exercise. Use the following:

Lemma (Spectral theorem for symmetric matrices)

Symmetric matrices are **orthogonally diagonalizable**, *i.e., for any symmetric matrix $A \in \mathbb{R}^{n\times n}$ there exist an orthogonal matrix $Q \in \mathbb{R}^{n\times n}$ and a diagonal matrix $D \in \mathbb{R}^{n\times n}$ s.t. $A = QDQ^{\mathrm{T}}$. The diagonal entries of D are the eigenvalues of A, and the column vectors of Q are eigenvectors of A. In particular, all eigenvalues of a symmetric matrix are real.*

Examples for definiteness

- $A := \begin{pmatrix} 1 & -2 \\ -2 & 4 \end{pmatrix} \succeq 0$ as $\Lambda(A) = \{0, 5\} \subseteq [0, \infty)$.

- $B := -A \preceq 0$ as $\Lambda(B) = \{-5, 0\} \subseteq (-\infty, 0]$.

- $C := \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix} \succ 0$ as $\Lambda(C) = \{1, 3\} \subseteq (0, \infty)$.

- $D := -C \prec 0$ as $\Lambda(D) = \{-3, -1\} \subseteq (-\infty, 0)$.

- $E := \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$ is neither positive semidefinite, nor negative semidefinite (we say E is **indefinite**), as $\Lambda(E) = \{-1, 3\}$.

More on positive definite matrices

- Let $A \in \mathbb{R}^{n\times n}$ be a symmetric positive definite matrix and let $X \in \mathbb{R}^{n\times r}$ with $n \geq r$ and $\mathrm{rk}(X) = r$. Then, the matrix $X^{\mathrm{T}}AX$ is symmetric positive definite (exercise).
- A useful criterion for checking positive definiteness:

Theorem (Sylvester's criterion for positive definiteness)

Let $A \in \mathbb{R}^{n\times n}$ be a symmetric matrix. Then,

$$A \succ 0 \quad \iff \quad \forall i \in \{1,\dots,n\} : \det(A_{1:i,1:i}) > 0.$$

The number $\det(A_{1:i,1:i})$ is called the i-th **leading principal minor** *of A. Therefore, a symmetric matrix is positive definite iff all of its leading principal minors are positive.*

$\Longrightarrow$ **Any symmetric positive definite matrix has a LU factorization!**
Even better: *We can factorize a symmetric positive definite matrix twice as quickly into triangular factors as a general matrix.*

Cholesky factorization

Definition (Cholesky factorization)

Let $A \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix. If there exists an upper-triangular matrix $R \in \mathbb{R}^{n \times n}$ with **positive diagonal entries** s.t.

$$A = R^{\mathrm{T}} R,$$

then we call this a **Cholesky factorization** of A.

The following is the main result of this section:

Theorem (Existence and uniqueness of Cholesky factorization)

Every symmetric positive definite matrix $A \in \mathbb{R}^{n \times n}$ admits a unique Cholesky factorization.

So, let's prove this ...

Existence of Cholesky factorization: Symmetric Gauß

Consider a symmetric positive definite matrix $A \in \mathbb{R}^{n\times n}$. Write

$$A = \left(\begin{array}{c|c} a_{11} & w^{\mathrm{T}} \\ \hline w & B \end{array}\right) \in \mathbb{R}^{n\times n}$$

with $a_{11} \in \mathbb{R}$, $w \in \mathbb{R}^{n-1}$ and a symmetric matrix $B \in \mathbb{R}^{(n-1)\times(n-1)}$.

Note that

- $a_{11} = \det(A_{1:1,1:1}) > 0$,
- $B \succ 0$ since $\langle x, Bx\rangle = \left\langle \begin{pmatrix} 0 \\ x \end{pmatrix}, A \begin{pmatrix} 0 \\ x \end{pmatrix} \right\rangle > 0$ for all $x \in \mathbb{R}^{n-1}\backslash\{0\}$.

First step of **symmetric Gaussian elimination**:

$$L_1 A L_1^{\mathrm{T}} = \left(\begin{array}{c|c} 1 & 0_{1\times(n-1)} \\ \hline 0_{(n-1)\times 1} & B - \frac{ww^{\mathrm{T}}}{a_{11}} \end{array}\right) =: A_1 \quad \text{with} \quad L_1 := \left(\begin{array}{c|c} \frac{1}{\sqrt{a_{11}}} & 0_{1\times(n-1)} \\ \hline -\frac{w}{a_{11}} & I_{n-1} \end{array}\right)$$

which we can equivalently write as

$$A = R_1^{\mathrm{T}} A_1 R_1 \quad \text{with} \quad R_1 := (L_1^{-1})^{\mathrm{T}} = \left(\begin{array}{c|c} \sqrt{a_{11}} & \frac{w^{\mathrm{T}}}{\sqrt{a_{11}}} \\ \hline 0_{(n-1)\times 1} & I_{n-1} \end{array}\right).$$

Recall $A_1 = L_1 A L_1^{\mathrm{T}} = \left(\begin{array}{c|c} 1 & 0_{1\times(n-1)} \\ \hline 0_{(n-1)\times 1} & B - \frac{ww^{\mathrm{T}}}{a_{11}} \end{array}\right)$.

Note that

- A_1 is symmetric,
- $A_1 = (L_1^{\mathrm{T}})^{\mathrm{T}} A (L_1^{\mathrm{T}}) \succ 0$ since $L_1^{\mathrm{T}} \in \mathbb{R}^{n\times n}$ is of full rank.

Therefore, we also have that the sub-matrix $B - \frac{ww^{\mathrm{T}}}{a_{11}} \in \mathbb{R}^{(n-1)\times(n-1)}$ is symmetric positive definite (same argument as when we deduced $B \succ 0$ from $A \succ 0$) and in particular, the $(1,1)$-entry of $B - \frac{ww^{\mathrm{T}}}{a_{11}}$ is positive.

$\Longrightarrow$ We can factor

$$A_1 = R_2^{\mathrm{T}} A_2 R_2$$

with $R_2 \in \mathbb{R}^{n\times n}$ upper-triangular with positive diagonal entries and A_2 being of the form $A_2 = \left(\begin{array}{c|c} I_2 & 0_{2\times(n-2)} \\ \hline 0_{(n-2)\times 2} & C \end{array}\right)$, using the same procedure as before applied to $B - \frac{ww^{\mathrm{T}}}{a_{11}}$. Then, again, the sub-matrix C is symmetric positive definite, and we can continue this process ...

$\Longrightarrow$ continue this process until we arrive at a factorization

$$A = (R_1^{\mathrm{T}} R_2^{\mathrm{T}} \cdots R_n^{\mathrm{T}}) I_n (R_n \cdots R_2 R_1) = R^{\mathrm{T}} R$$

with $R := R_n \cdots R_2 R_1 \in \mathbb{R}^{n \times n}$ upper-triangular and having positive diagonal entries. This is a Cholesky factorization of A!

Next: Remains to show uniqueness.

Theorem (Existence and uniqueness of Cholesky factorization)

Every symmetric positive definite matrix $A \in \mathbb{R}^{n \times n}$ admits a unique Cholesky factorization.

Proof: Symmetric Gaussian elimination provides existence of a Cholesky factorization (argument can be made rigorous via induction).

For uniqueness, suppose that $R, M \in \mathbb{R}^{n \times n}$ are two upper-triangular matrices with positive diagonal entries such that

$$A = R^{\mathrm{T}} R = M^{\mathrm{T}} M.$$

Note that $D := M R^{-1}$ is an upper-triangular matrix, but also, since

$$D = M R^{-1} = (M^{\mathrm{T}})^{-1} R^{\mathrm{T}} = (D^{-1})^{\mathrm{T}},$$

it must be lower-triangular as well, hence diagonal.

Noting that $I_n = D^{\mathrm{T}} D = D^2$, the diagonal entries of D are all ± 1. Finally, since $DR = M$ and the diagonal entries of R and M are positive, we must have that $R = M$. □

Example: Computing the Cholesky factorization

Consider the symmetric positive definite matrix

$$A := \begin{pmatrix} 16 & -8 & 12 \\ -8 & 5 & -9 \\ 12 & -9 & 22 \end{pmatrix} \in \mathbb{R}^{3\times 3}.$$

We illustrate symmetric Gaussian elimination:

L_1: Eliminate the sub-diagonal entries in the first column of A by adding $\frac{1}{2}$/-$\frac{3}{4}$ times row 1 to row 2/3, and multiply the first row by $\frac{1}{\sqrt{a_{11}}} = \frac{1}{4}$:

$$L_1 A = \begin{pmatrix} 4 & -2 & 3 \\ 0 & 1 & -3 \\ 0 & -3 & 13 \end{pmatrix} \quad \text{with} \quad L_1 := \begin{pmatrix} 1/4 & 0 & 0 \\ 1/2 & 1 & 0 \\ -3/4 & 0 & 1 \end{pmatrix}.$$

Next, we right-multiply $L_1 A$ with L_1^{T} which creates a 1 in the $(1,1)$ entry and zeros in the $(1,2)$ and $(1,3)$ entries:

$$L_1 A L_1^{\mathrm{T}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -3 \\ 0 & -3 & 13 \end{pmatrix}.$$

Recall $L_1 A L_1^{\mathrm{T}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -3 \\ 0 & -3 & 13 \end{pmatrix}$.

L_2: Eliminate sub-diagonal entry in second column of $L_1 A L_1^{\mathrm{T}}$ by adding 3 times row 2 to row 3 (and multiply the second row by $\frac{1}{\sqrt{1}} = 1$):

$$L_2 L_1 A L_1^{\mathrm{T}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -3 \\ 0 & 0 & 4 \end{pmatrix} \quad \text{with} \quad L_2 := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 3 & 1 \end{pmatrix}.$$

Next, we right-multiply $L_2 L_1 A L_1^{\mathrm{T}}$ with L_2^{T} which creates a zero in the $(2, 3)$ entry:

$$L_2 L_1 A L_1^{\mathrm{T}} L_2^{\mathrm{T}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 4 \end{pmatrix}.$$

Recall $L_2 L_1 A L_1^{\mathrm{T}} L_2^{\mathrm{T}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 4 \end{pmatrix}$.

L_3: We multiply the third row of $L_2 L_1 A L_1^{\mathrm{T}} L_2^{\mathrm{T}}$ by $\frac{1}{\sqrt{4}} = \frac{1}{2}$:

$$L_3 L_2 L_1 A L_1^{\mathrm{T}} L_2^{\mathrm{T}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix} \quad \text{with} \quad L_3 := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{1}{2} \end{pmatrix}.$$

Finally, we right-multiply $L_3 L_2 L_1 A L_1^{\mathrm{T}} L_2^{\mathrm{T}}$ by L_3^{T} which creates a 1 in the $(3,3)$ entry:

$$L_3 L_2 L_1 A L_1^{\mathrm{T}} L_2^{\mathrm{T}} L_3^{\mathrm{T}} = I_3.$$

We find that $A = R^{\mathrm{T}} R$ with

$$R := [L_1^{-1} L_2^{-1} L_3^{-1}]^{\mathrm{T}} = \left[\begin{pmatrix} 4 & 0 & 0 \\ -2 & 1 & 0 \\ 3 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -3 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix} \right]^{\mathrm{T}} = \begin{pmatrix} 4 & -2 & 3 \\ 0 & 1 & -3 \\ 0 & 0 & 2 \end{pmatrix}$$

is the unique Cholesky factorization of A.

Cholesky factorization via symmetric Gauß: Algorithm

To obtain the Cholesky factorization $A = R^{\mathrm{T}} R$ of a given symmetric positive definite matrix $A \in \mathbb{R}^{n \times n}$, do as follows:

$R = A$
for $i = 1, \ldots, n$ **do**
  for $j = i+1, \ldots, n$ **do**
    $R_{j,j:n} = R_{j,j:n} - \frac{R_{i,j:n} R_{ij}}{R_{ii}}$
  end for
  $R_{i,i:n} = \frac{R_{i,i:n}}{\sqrt{R_{ii}}}$
end for.

Theorem

The above algorithm requires $\sim \frac{1}{3} n^3$ *flops.*

This is only half the cost of Gaussian elimination!

Solving linear systems via Cholesky factorization

For a given symmetric positive definite matrix $A \in \mathbb{R}^{n \times n}$ and a vector $b \in \mathbb{R}^n$, consider the problem of finding $x \in \mathbb{R}^n$ such that $Ax = b$.

The standard way to solve the system in this case is by Cholesky factorization: If $A = R^{\mathrm{T}} R$ is the Cholesky factorization of A, we have

$$Ax = b \quad \Longleftrightarrow \quad R^{\mathrm{T}} R x = b \quad \Longleftrightarrow \quad \begin{cases} R^{\mathrm{T}} y = b, \\ Rx = y. \end{cases}$$

Therefore, once the Cholesky factorization is computed ($\mathcal{O}(n^3)$ flops), we can first solve $R^{\mathrm{T}} y = b$ for y by forward substitution ($\mathcal{O}(n^2)$ flops) and then $Rx = y$ for x by backward substitution ($\mathcal{O}(n^2)$ flops).

4.5 Least Squares Problems

Over-determined linear systems

Given $A = (a_{ij}) \in \mathbb{R}^{m \times n}$ with $m > n$, and $b = (b_1, \ldots, b_m)^{\mathrm{T}} \in \mathbb{R}^m$.

Problem: Find $x = (x_1, \ldots, x_n)^{\mathrm{T}} \in \mathbb{R}^n$ s.t.

$$Ax = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{n1} & \cdots & a_{nn} \\ a_{n+1,1} & \cdots & a_{n+1,n} \\ \vdots & & \vdots \\ a_{m1} & \cdots & a_{mn} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = \begin{pmatrix} b_1 \\ \vdots \\ b_n \\ b_{n+1} \\ \vdots \\ b_m \end{pmatrix} = b.$$

Such a problem does not admit a solution in general: consider e.g.,

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

The least squares problem

Observation: Given $A \in \mathbb{R}^{m \times n}$, $m > n$, and $b \in \mathbb{R}^m$,

$$[\, \exists x \in \mathbb{R}^n : Ax = b \,] \quad \Longleftrightarrow \quad b \in \mathscr{R}(A).$$

Noting that $\dim(\mathscr{R}(A)) \leq n < m = \dim(\mathbb{R}^m)$, such an over-determined system $Ax = b$ is only solvable for special choices of $b \in \mathbb{R}^m$.

$\Longrightarrow$ Consider the following generalized problem:

Find $\quad x \in \mathbb{R}^n \quad$ s.t. $\quad r := Ax - b \quad$ is as small as possible.

We call r the **residual**. To measure the size of r, use the Euclidean norm.

Definition (Least squares problem)

Given $A \in \mathbb{R}^{m \times n}$, $m \geq n$, and $b \in \mathbb{R}^m$, we call the following problem the **least squares problem** corresponding to the matrix A and the vector b:

Minimize $\|Av - b\|_2$ over $v \in \mathbb{R}^n$.

A vector $x \in \mathbb{R}^n$ is called a solution to the least squares problem iff

$$\|Ax - b\|_2 = \inf_{v \in \mathbb{R}^n} \|Av - b\|_2.$$

Motivation: Interpolation vs. least squares fitting

Suppose we are given data points $(t_1, y_1), \dots, (t_n, y_n)$ with $t_1, \dots, t_n \in \mathbb{R}$ distinct and $y_1, \dots, y_n \in \mathbb{R}$.

(i) **Polynomial interpolation:**
There exists a unique polynomial $p(t) = \sum_{k=0}^{n-1} p_k t^k$ of degree $n-1$ such that $p(t_i) = y_i$ for all $i \in \{1, \dots, n\}$. (**polynomial interpolant**)

The coefficients $p_0, \dots, p_{n-1} \in \mathbb{R}$ are uniquely determined from

$$V \begin{pmatrix} p_0 \\ \vdots \\ p_{n-1} \end{pmatrix} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}, \qquad V := \begin{pmatrix} 1 & t_1 & t_1^2 & \cdots & t_1^{n-1} \\ 1 & t_2 & t_2^2 & \cdots & t_2^{n-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & t_n & t_n^2 & \cdots & t_n^{n-1} \end{pmatrix} \in \mathbb{R}^{n \times n}.$$

Note that the so-called **Vandermonde matrix** $V = V(t_1, \dots, t_n)$ is invertible since the values $\{t_i\}$ are distinct.
$\Longrightarrow$ Great! Or not? ...
Drawback: Large oscillations near the ends of the interval $[t_1, t_n]$.

Motivation: Interpolation vs. least squares fitting

Data $(t_1, y_1), \ldots, (t_n, y_n)$ with $t_1, \ldots, t_n \in \mathbb{R}$ distinct, $y_1, \ldots, y_n \in \mathbb{R}$.

(ii) **Least squares fitting:** Ansatz: $p(t) = \sum_{k=0}^{N-1} p_k t^k$ with $N < n$.

The condition $p(t_i) = y_i$ for $i \in \{1, \ldots, n\}$ leads to

$$Ap_{\text{coeff}} := \begin{pmatrix} 1 & t_1 & t_1^2 & \cdots & t_1^{N-1} \\ 1 & t_2 & t_2^2 & \cdots & t_2^{N-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & t_n & t_n^2 & \cdots & t_n^{N-1} \end{pmatrix} \begin{pmatrix} p_0 \\ \vdots \\ p_{N-1} \end{pmatrix} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} =: b.$$

which may not have a solution. Instead, we choose the coefficient vector $p_{\text{coeff}} = (p_0, \ldots, p_{N-1})^{\mathrm{T}} \in \mathbb{R}^N$ s.t.

$$\|Ap_{\text{coeff}} - b\|_2 = \inf_{v \in \mathbb{R}^N} \|Av - b\|_2.$$

The least squares fit $p(t) = \sum_{k=0}^{N-1} p_k t^k$ minimizes the quantity $\sqrt{\sum_{i=1}^{n} |p(t_i) - y_i|^2}$ among polynomials of degree at most $N - 1$.
$\Longrightarrow$ **The least squares soln does not interpolate the data points, but it describes the overall behavior better than the interpolant.**

Example: Data points

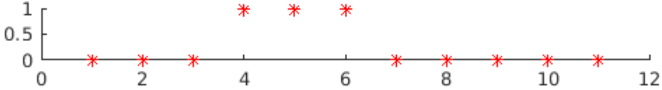

Example: Interpolant (polynomial of degree 10)

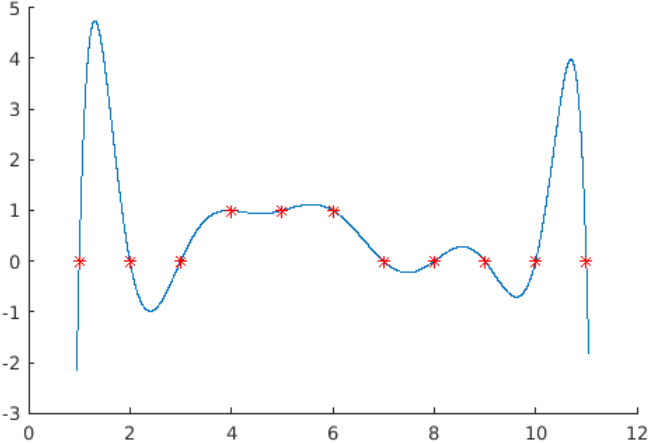

Example: Degree 7 polynomial least squares fit

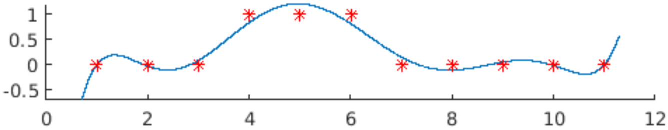

The main questions

Recall:

Definition (Least squares problem)

Given $A \in \mathbb{R}^{m\times n}$, $m \geq n$, and $b \in \mathbb{R}^m$, we call the following problem the **least squares problem** corresponding to the matrix A and the vector b:

$$\text{Minimize } \|Av - b\|_2 \text{ over } v \in \mathbb{R}^n.$$

A vector $x \in \mathbb{R}^n$ is called a solution to the least squares problem iff

$$\|Ax - b\|_2 = \inf_{v\in\mathbb{R}^n} \|Av - b\|_2.$$

- **Existence**: Is there a soln to the LS problem for any choices of $A \in \mathbb{R}^{m\times n}$ with $m \geq n$ and $b \in \mathbb{R}^m$?
- **Uniqueness**: If there exists a soln to the LS problem, is this unique?
- **Computation**: If there exist solns to the LS problems, how can we find them?

An "equivalent" minimization problem

Recall LS problem:

$$\text{Minimize } \|Av - b\|_2 \text{ over } v \in \mathbb{R}^n.$$

If there exists $x \in \mathbb{R}^n$ s.t. $\|Ax - b\|_2 = \inf_{v \in \mathbb{R}^n} \|Av - b\|_2$, then we call this minimizer x a solution to the LS problem.

Introduce an *"equivalent" minimization problem*:

$$\text{Minimize } \|w - b\|_2 \text{ over } w \in \mathscr{R}(A). \tag{1}$$

If there exists $y \in \mathscr{R}(A)$ s.t. $\|y - b\|_2 = \inf_{w \in \mathscr{R}(A)} \|w - b\|_2$, then we call this minimizer y a solution to the above minimization problem.

Observations:

- If $\exists$ solution $x \in \mathbb{R}^n$ to LS, then $y = Ax \in \mathscr{R}(A)$ is a solution to (1).
- If $\exists$ solution $y \in \mathscr{R}(A)$ to (1), then any $x \in \mathbb{R}^n$ satisfying $Ax = y$ is a solution to LS.
- There holds $\inf_{v \in \mathbb{R}^n} \|Av - b\|_2 = \inf_{w \in \mathscr{R}(A)} \|w - b\|_2$.

Geometric illustration of the LS problem

Ingredients for existence proof

Theorem (Existence of solutions to the normal equation)

Let $A \in \mathbb{R}^{m\times n}$. Then, for any $b \in \mathbb{R}^m$ there exists a solution $x \in \mathbb{R}^n$ to the **normal equation** $A^{\mathrm{T}}Ax = A^{\mathrm{T}}b$.

Proof: Need to show that $A^{\mathrm{T}}b \in \mathscr{R}(A^{\mathrm{T}}A)$ for any $b \in \mathbb{R}^m$. We are going to show that $\mathscr{R}(A^{\mathrm{T}}) = \mathscr{R}(A^{\mathrm{T}}A)$:

$$\mathscr{R}(A^{\mathrm{T}}) = [\mathscr{N}(A)]^{\perp} = [\mathscr{N}(A^{\mathrm{T}}A)]^{\perp} = \mathscr{R}((A^{\mathrm{T}}A)^{\mathrm{T}}) = \mathscr{R}(A^{\mathrm{T}}A).$$

(Used $\mathscr{N}(A) = \mathscr{N}(A^{\mathrm{T}}A)$ and $[\mathscr{N}(M)]^{\perp} = \mathscr{R}(M^{\mathrm{T}})$ for any matrix M.) $\square$

Theorem (Orthogonal projector onto range of matrix)

Let $A \in \mathbb{R}^{m\times n}$. Then,

(i) $\mathscr{R}(A)$ *and* $\mathscr{N}(A^{\mathrm{T}})$ *are complementary subspaces of* $\mathbb{R}^m$,

(ii) $\mathscr{R}(A) \perp \mathscr{N}(A^{\mathrm{T}})$.

In particular, $\exists$ a unique projector $P \in \mathbb{R}^{m\times m}$ s.t. $\mathscr{R}(P) = \mathscr{R}(A)$ and $\mathscr{N}(P) = \mathscr{N}(A^{\mathrm{T}})$, and P is the unique orthogonal projector onto $\mathscr{R}(A)$.

Existence and Uniqueness results

Theorem (Existence and uniqueness result for least squares problems)

Let $A \in \mathbb{R}^{m \times n}$, $m \geq n$, and $b \in \mathbb{R}^m$. Let $P \in \mathbb{R}^{m \times m}$ be the orthogonal projector onto $\mathscr{R}(A)$. Then, we have the following:

(i) *$\exists$ a unique solution to the minimization problem* (1)*, i.e., a unique $y \in \mathscr{R}(A)$ with $\|y - b\|_2 = \inf_{w \in \mathscr{R}(A)} \|w - b\|_2$. This soln is given by*

$$y = Pb.$$

(ii) *$\exists$ a solution to the least squares problem, i.e., $x \in \mathbb{R}^n$ satisfying $\|Ax - b\|_2 = \inf_{v \in \mathbb{R}^n} \|Av - b\|_2$. Moreover, $x \in \mathbb{R}^n$ is a solution iff*

$$Ax = Pb, \text{ or equivalently, } A^{\mathrm{T}}Ax = A^{\mathrm{T}}b.$$

(iii) *The least squares problem has a unique solution iff $\mathrm{rk}(A) = n$.*

Proof of (i)

We need to show that the minimization problem

$$\text{Minimize } \|w-b\|_2 \text{ over } w \in \mathscr{R}(A)$$

has the unique solution $y = Pb \in \mathscr{R}(A)$. (Note $Pb \in \mathscr{R}(P) = \mathscr{R}(A)$.)

We have for any $w \in \mathscr{R}(A) \backslash \{Pb\}$ that

$$\|w-b\|_2^2 = \|(w-Pb)+(Pb-b)\|_2^2 = \|w-Pb\|_2^2 + \|Pb-b\|_2^2 > \|Pb-b\|_2^2,$$

where we have used that $\langle \underbrace{w-Pb}_{\in \mathscr{R}(P)}, \underbrace{Pb-b}_{\in \mathscr{N}(P)} \rangle = 0$.

$\Longrightarrow y = Pb$ is the unique element in $\mathscr{R}(A)$ satisfying

$$\|y-b\|_2 = \inf_{w \in \mathscr{R}(A)} \|w-b\|_2.$$

Proof of (ii)

Need to show the following: $\exists\, x \in \mathbb{R}^n : \|Ax - b\|_2 = \inf_{v \in \mathbb{R}^n} \|Av - b\|_2$, and that $x \in \mathbb{R}^n$ is a solution iff $Ax = Pb$ iff $A^{\mathrm{T}}Ax = A^{\mathrm{T}}b$.

By (i), any $x \in \mathbb{R}^n$ satisfying $Ax = Pb$ is a solution to LS. Conversely, if $x \in \mathbb{R}^n$ is a solution to LS, then $y = Ax$ is a solution to $\|y - b\|_2 = \inf_{w \in \mathscr{R}(A)} \|w - b\|_2$ and consequently, $Ax = Pb$.

Remains to show that for $x \in \mathbb{R}^n$, we have $Ax = Pb \Longleftrightarrow A^{\mathrm{T}}Ax = A^{\mathrm{T}}b$.

"$\Longrightarrow$" Let $x \in \mathbb{R}^n$ with $Ax = Pb$. Then,

$$Ax - b = Pb - b \in \mathcal{N}(P) = \mathcal{N}(A^{\mathrm{T}}) \qquad \Longrightarrow \qquad A^{\mathrm{T}}Ax = A^{\mathrm{T}}b.$$

"$\Longleftarrow$" Let $x \in \mathbb{R}^n$ with $A^{\mathrm{T}}Ax = A^{\mathrm{T}}b$. Then $Ax - b \in \mathcal{N}(A^{\mathrm{T}}) = \mathcal{N}(P)$ and hence,

$$Ax - Pb = (I_m - P)Ax + P(Ax - b) = 0,$$

where we have used that $Ax \in \mathscr{R}(A) = \mathscr{R}(P) = \mathcal{N}(I_m - P)$.

Proof of (iii)

Need to show: Solution to LS unique iff A has full rank.

By (ii), LS has a unique soln iff $A^{\mathrm{T}}Ax = A^{\mathrm{T}}b$ has a unique soln $x \in \mathbb{R}^n$, i.e., iff $A^{\mathrm{T}}A \in \mathbb{R}^{n \times n}$ is invertible, i.e., iff $\mathrm{rk}(A^{\mathrm{T}}A) = n$, i.e., iff $\mathrm{rk}(A) = n$.

(Recall $\mathscr{R}(A^{\mathrm{T}}A) = \mathscr{R}(A^{\mathrm{T}}) \implies \mathrm{rk}(A^{\mathrm{T}}A) = \mathrm{rk}(A^{\mathrm{T}}) = \mathrm{rk}(A)$.)

□

Solution of the full-rank least squares problem

Let $A \in \mathbb{R}^{m \times n}$, $m \geq n$, and assume that $\mathrm{rk}(A) = n$. Then, the unique solution to LS is given by

$$A^{\mathrm{T}} A x = A^{\mathrm{T}} b \quad \Longrightarrow \quad x = (A^{\mathrm{T}} A)^{-1} A^{\mathrm{T}} b.$$

We find that

$$x = A^{\dagger} b \in \mathbb{R}^n, \quad \text{where} \quad A^{\dagger} := (A^{\mathrm{T}} A)^{-1} A^{\mathrm{T}} \in \mathbb{R}^{n \times m}.$$

The matrix $A^{\dagger}$ is the **Moore–Penrose inverse** (or **pseudoinverse**) of A.

The Moore–Penrose inverse is a generalization of the matrix inverse and is being discussed extensively on the problem sheets.

Solution Algorithm 1: via normal eqn & Cholesky

Let $A \in \mathbb{R}^{m\times n}$, $m \geq n$, $b \in \mathbb{R}^m$ and assume $\mathrm{rk}(A) = n$. Then,

$$A^{\mathrm{T}}A \in \mathbb{R}^{n\times n} \quad \text{is symmetric positive definite.}$$

Indeed, we have $(A^{\mathrm{T}}A)^{\mathrm{T}} = A^{\mathrm{T}}A$ and

$$\langle x, A^{\mathrm{T}}Ax\rangle = \langle Ax, Ax\rangle = \|Ax\|_2^2 > 0 \qquad \forall x \in \mathbb{R}^n\backslash\{0\}.$$

Here, we have used that $Ax \in \mathbb{R}^m\backslash\{0\}$ for $x \in \mathbb{R}^n\backslash\{0\}$ since $\mathrm{rk}(A) = n$.

Therefore, $A^{\mathrm{T}}A$ has a unique Cholesky factn $A^{\mathrm{T}}A = R^{\mathrm{T}}R$ and we have

$$A^{\mathrm{T}}Ax = A^{\mathrm{T}}b \quad \Longleftrightarrow \quad R^{\mathrm{T}}Rx = A^{\mathrm{T}}b.$$

Algorithm:

1) Compute $\tilde{A} := A^{\mathrm{T}}A \in \mathbb{R}^{n\times n}$ and $\tilde{b} := A^{\mathrm{T}}b \in \mathbb{R}^n$.
2) Compute the Cholesky factorization $\tilde{A} = R^{\mathrm{T}}R$ of $\tilde{A}$.
3) Solve the lower-triangular system $R^{\mathrm{T}}z = \tilde{b}$ for $z \in \mathbb{R}^n$.
4) Solve the upper-triangular system $Rx = z$ for $x \in \mathbb{R}^n$.

1) Compute $\tilde{A} := A^{\mathrm{T}} A \in \mathbb{R}^{n \times n}$ and $\tilde{b} := A^{\mathrm{T}} b \in \mathbb{R}^{n}$.
2) Compute the Cholesky factorization $\tilde{A} = R^{\mathrm{T}} R$ of $\tilde{A}$.
3) Solve the lower-triangular system $R^{\mathrm{T}} z = \tilde{b}$ for $z \in \mathbb{R}^{n}$.
4) Solve the upper-triangular system $Rx = z$ for $x \in \mathbb{R}^{n}$.

Theorem

This algorithm requires $\sim mn^2 + \frac{1}{3}n^3$ *flops.*

Solution Algorithm 2: via reduced QR

Let $A \in \mathbb{R}^{m\times n}$, $m \geq n$, $b \in \mathbb{R}^m$, and assume $A = \hat{Q}\hat{R}$ reduced QR factn.

Then, $x \in \mathbb{R}^n$ is soln to LS iff $A^{\mathrm{T}}Ax = A^{\mathrm{T}}b$ iff

$$\hat{R}^{\mathrm{T}}\hat{Q}^{\mathrm{T}}\hat{Q}\hat{R}x = \hat{R}^{\mathrm{T}}\hat{Q}^{\mathrm{T}}b \quad \Longrightarrow \quad \hat{R}^{\mathrm{T}}\hat{R}x = \hat{R}^{\mathrm{T}}\hat{Q}^{\mathrm{T}}b.$$

Observe: If A is of full rank, then $\hat{R}$ is invertible and thus,

$$\hat{R}x = \hat{Q}^{\mathrm{T}}b.$$

Assume $\mathrm{rk}(A) = n$. Do the following:

1) Compute a reduced QR factorization $A = \hat{Q}\hat{R}$ of A.
2) Compute $\tilde{b} = \hat{Q}^{\mathrm{T}}b \in \mathbb{R}^n$.
3) Solve the upper-triangular system $\hat{R}x = \tilde{b}$ for $x \in \mathbb{R}^n$.

Using Householder find:

Theorem

This algorithm requires $\sim 2mn^2 - \frac{2}{3}n^3$ flops.

Solution Algorithm 3: via reduced SVD

Let $A \in \mathbb{R}^{m\times n}$, $m \geq n$, $b \in \mathbb{R}^m$, and assume $A = \hat{U}\hat{\Sigma}V^{\mathrm{T}}$ reduced SVD.

Then, $x \in \mathbb{R}^n$ is a solution to LS iff $A^{\mathrm{T}}Ax = A^{\mathrm{T}}b$ iff

$$V\hat{\Sigma}^{\mathrm{T}}\hat{U}^{\mathrm{T}}\hat{U}\hat{\Sigma}V^{\mathrm{T}}x = V\hat{\Sigma}^{\mathrm{T}}\hat{U}^{\mathrm{T}}b \quad\Longrightarrow\quad V\hat{\Sigma}^{\mathrm{T}}\hat{\Sigma}V^{\mathrm{T}}x = V\hat{\Sigma}^{\mathrm{T}}\hat{U}^{\mathrm{T}}b.$$

Observe: If A is of full rank, then $V\hat{\Sigma}^{\mathrm{T}} \in \mathbb{R}^{n\times n}$ is invertible and thus,

$$\hat{\Sigma}V^{\mathrm{T}}x = \hat{U}^{\mathrm{T}}b.$$

Assume $\mathrm{rk}(A) = n$. Do the following:

1) Compute a reduced SVD $A = \hat{U}\hat{\Sigma}V^{\mathrm{T}}$ of A.
2) Compute $\tilde{b} = \hat{U}^{\mathrm{T}}b \in \mathbb{R}^n$.
3) Solve the diagonal system $\hat{\Sigma}z = \tilde{b}$ for $z \in \mathbb{R}^n$.
4) Compute $x = Vz \in \mathbb{R}^n$.

Theorem

This algorithm requires $\sim 2mn^2 + 11n^3$ *flops.*

End of “Chapter 4: Linear Systems and Least Squares Problems”.