

Elementary Matrices

This turns row operations into matrix multiplication.

It lets us use matrix algebra to reason about elimination.

Def: An elementary matrix is a matrix obtained from the identity matrix by doing one row operation.

Eg:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \xrightarrow{R_2 \leftarrow 2R_2} \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \xrightarrow{R_3 \leftarrow 3} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \xrightarrow{R_1 \leftrightarrow R_2} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Fact: If E is the elementary matrix for a row operation, then

$$E \cdot A = \left(\text{the matrix obtained by doing the same row operation to } A \right)$$



Eg:

$$A = \begin{pmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 5 & -3 & -6 & -6 \end{pmatrix} \quad R_3 \leftarrow 5R_1$$

$$\boxed{\begin{pmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 0 & -8 & -1 & -21 \end{pmatrix}}$$

$$E = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad R_3 \leftarrow 5R_1$$

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -5 & 0 & 1 \end{pmatrix} \quad \text{|| same!}$$

$$EA = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -5 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 5 & -3 & -6 & -6 \end{pmatrix} = \boxed{\begin{pmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 0 & -8 & -1 & -21 \end{pmatrix}}$$

Left-multiplication by $\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -5 & 0 & 1 \end{pmatrix}$ does $R_3 \leftarrow 5R_1$

Fact: An elementary matrix is invertible. Its inverse is the elementary matrix that un-does the row operation.

Why?

$$E_1 = \begin{pmatrix} \text{elementary} \\ \text{matrix for} \\ R_3 \leftarrow 5R_1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -5 & 0 & 1 \end{pmatrix}$$

$$E_2 = \begin{pmatrix} \text{elementary} \\ \text{matrix for} \\ R_3 \leftarrow 5R_1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 5 & 0 & 1 \end{pmatrix}$$

$$E_2 E_1 = E_2 E_1 I_3 = E_2 (E_1 I_3) = \begin{pmatrix} \text{first multiply } I_3 \text{ by } E_1 \\ \text{then multiply that by } E_2 \end{pmatrix}$$

What does that do?

I_3

$R_3 \leftarrow 5R_1$

mult. by E_1

$E_1 I_3$

$R_3 \leftarrow 5R_1$

mult. by E_2

$E_2 (E_1 I_3)$

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

$R_3 \leftarrow 5R_1$

mult. by E_1

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -5 & 0 & 1 \end{pmatrix}$$

$R_3 \leftarrow 5R_1$

mult. by E_2

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

This does the row operation $R_3 \leftarrow 5R_1$, then un-does it!
So you get I_3 back.

$$E_2 E_1 = I_3 \text{ means } E_1^{-1} = E_2$$

This illustrates the following important point:

If E is an elementary matrix, then
compute $E \cdot A$ by doing a row operation.
Not by multiplying matrices!

What if you do multiple row operations?

Consider these elementary matrices:

$$E_1: R_1 \leftarrow 2R_2 \quad E_2: R_2 \leftarrow 2 \quad E_3: R_2 \leftrightarrow R_3$$

Let's apply these in order to a matrix A :

$$A \xrightarrow{R_1 \leftarrow 2R_2} E_1 A \xrightarrow{R_2 \leftarrow 2} \underset{\substack{\text{do } R_2 \leftarrow 2 \\ \text{to } E_1 A}}{E_2 (E_1 A)} \xrightarrow{R_2 \leftrightarrow R_3} \underset{\substack{\text{do } R_2 \leftrightarrow R_3 \\ \text{to } E_2 E_1 A}}{E_3 (E_2 (E_1 A))} = (E_3 E_2 E_1) A$$

Why did the elementary matrices end up in the opposite order?

$$E_3 E_2 E_1 A = E_3 E_2 (E_1 A)$$

= first multiply by E_1 , then by E_2 , then E_3

Application to Invertibility:

Suppose $A \xrightarrow{\text{RREF}} I_n$. That means you can do some number of row operations to A to get I_n . Let E_0, \dots, E_r be the elementary matrices for these row operations. Then

$$I_n = (E_r \cdots E_2 E_1) A$$

$$\Rightarrow A^{-1} = E_r \cdots E_2 E_1$$

In particular, A is invertible!

It also tells us how to compute A^{-1} using row operations:

$$\begin{array}{l} \xrightarrow{\text{column-first matrix multiplication}} \\ \Rightarrow C(A|B) = (CA|CB) \end{array}$$

$$\begin{aligned} (E_r \cdots E_2 E_1) (A | I_n) &= ((E_r \cdots E_2 E_1) A | (E_r \cdots E_2 E_1) I_n) \\ &= (I_n | A^{-1}) \end{aligned}$$

This means if you do the same row operations to I_n , then you get A^{-1} : that's our algorithm from last time.

Triangular Matrices

Def: A matrix is **upper/lower triangular** if all of the entries below/above the main diagonal are zero.

upper triangular

$$\begin{pmatrix} \text{green circle} & \text{green circle} & \text{green circle} & \text{green circle} \\ 0 & \text{green circle} & \text{green circle} & \text{green circle} \\ 0 & 0 & \text{green circle} & \text{green circle} \end{pmatrix}$$

main diagonal

lower triangular

$$\begin{pmatrix} \text{green circle} & 0 & 0 & 0 \\ \text{green circle} & 0 & 0 & 0 \\ \text{green circle} & 0 & 0 & 0 \end{pmatrix}$$

main diagonal

green circle = any number

Def: A matrix is **upper/lower unitriangular** if it is upper/lower triangular and all diagonal entries = 1.

upper unitriangular

$$\begin{pmatrix} 1 & \text{green circle} & \text{green circle} & \text{green circle} \\ 0 & 1 & \text{green circle} & \text{green circle} \\ 0 & 0 & 1 & \text{green circle} \end{pmatrix}$$

lower unitriangular

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ \text{green circle} & 1 & 0 & 0 \\ \text{green circle} & \text{green circle} & 1 & 0 \end{pmatrix}$$

green circle = any number

Eg: These matrices are upper-triangular but not unitriangular:

$$\begin{pmatrix} 1 & 4 & 5 & 7 \\ 0 & 2 & 6 & 8 \\ 0 & 0 & 3 & 9 \end{pmatrix} \quad \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & 4 & 5 & 7 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 9 \end{pmatrix} \quad \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

green circle = any number

(including zero!)

Eg: These matrices are lower uniΔular:

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

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

$$\begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 5 & 1 \\ 4 & 6 & 7 \end{pmatrix}$$

● = any number (including zero!)

NB: A matrix is **diagonal** \Leftrightarrow it is both upper- & lower-Δular: that means all off-diagonal entries are zero.

Eg: A matrix in REF is upper-Δular:

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

Eg: The elementary matrix for $R_i \leftarrow cR_i$, $i \geq j$

(add a multiple of a row to a row below it)

is lower-uniΔular:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \xrightarrow{R_3 \leftarrow 2R_1} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix}$$

Fact: If A & B are $n \times n$ upper (uni)Δular matrices then so are AB and A^{-1} (if A is invertible). Likewise for lower (uni)Δular matrices.

Eg: $\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 2 & 3 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 4 & 1 & 0 \\ 5 & 6 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 5 & 1 & 0 \\ 19 & 9 & 1 \end{pmatrix}$

$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 2 & 3 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} \text{matrix inversion} \\ \text{procedure...} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 1 & -3 & 1 \end{pmatrix}$

NB: Any square, **uni**lular matrix is invertible:

$$\begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix} \quad \bullet = \text{pivots}$$

LU Decompositions:

If Gaussian elimination on A requires no row swaps, then

$$A = LU$$

where:

L is lower **uni**lular

U is an REF for A (\Rightarrow upper **Δ**lular)

Using the prescribed algorithm, not your clever row ops!

Eg: $\begin{pmatrix} 2 & 1 & 0 \\ 4 & 4 & 4 \\ 6 & 1 & 0 \end{pmatrix} = A = LU = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & -1 & 1 \end{pmatrix} \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 0 & 0 & 4 \end{pmatrix}$

How does this speed up solving $Ax=b$?

How to solve $Ax=b$ using $A=LU$:

(1) Solve $Ly=b$ using substitution.

(2) Solve $Ux=y$ using substitution.

Then $Ax = (LU)x = L(Ux) = Ly = b$ ✓

Eg: Solve $\begin{pmatrix} 2 & 1 & 0 \\ 4 & 4 & 4 \\ 6 & 1 & 0 \end{pmatrix}x = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$ using

$$\begin{pmatrix} 2 & 1 & 0 \\ 4 & 4 & 4 \\ 6 & 1 & 0 \end{pmatrix} = A = LU = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & -1 & 1 \end{pmatrix} \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 0 & 0 & 4 \end{pmatrix}$$

(1) Solve $\begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & -1 & 1 \end{pmatrix}y = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$ using substitution.

$$\begin{pmatrix} 1 & 0 & 0 & | & 1 \\ 2 & 1 & 0 & | & 0 \\ 3 & -1 & 1 & | & 1 \end{pmatrix} \xrightarrow{\substack{R_2 \leftarrow 2R_1 \\ R_3 \leftarrow 3R_1}} \begin{pmatrix} 1 & 0 & 0 & | & 1 \\ 0 & 1 & 0 & | & -2 \\ 0 & -1 & 1 & | & -2 \end{pmatrix}$$

$$\xrightarrow{R_3 \leftarrow R_2} \begin{pmatrix} 1 & 0 & 0 & | & 1 \\ 0 & 1 & 0 & | & -2 \\ 0 & 0 & 1 & | & -4 \end{pmatrix} \Rightarrow y = \begin{pmatrix} 1 \\ -2 \\ -4 \end{pmatrix}$$

(2) Solve $\begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 0 & 0 & 4 \end{pmatrix}x = \begin{pmatrix} 1 \\ -2 \\ -4 \end{pmatrix}$ using substitution.

$$\begin{pmatrix} 2 & 1 & 0 & | & 1 \\ 0 & 2 & 4 & | & -2 \\ 0 & 0 & 4 & | & -4 \end{pmatrix} \xrightarrow{R_3 \div 4} \begin{pmatrix} 2 & 1 & 0 & | & 1 \\ 0 & 2 & 4 & | & -2 \\ 0 & 0 & 1 & | & -1 \end{pmatrix}$$

$$R_2 \leftarrow 4R_3 \rightarrow \left(\begin{array}{ccc|c} 2 & 1 & 0 & 1 \\ 0 & 2 & 0 & 2 \\ 0 & 0 & 1 & -1 \end{array} \right) \quad R_2 \div 2 \rightarrow \left(\begin{array}{ccc|c} 2 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & -1 \end{array} \right)$$

$$R_1 \leftarrow R_2 \rightarrow \left(\begin{array}{ccc|c} 2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & -1 \end{array} \right) \quad R_1 \div 2 \rightarrow \left(\begin{array}{ccc|c} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & -1 \end{array} \right)$$

$$\rightsquigarrow x = \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}$$

Check: $\left(\begin{array}{ccc} 2 & 1 & 0 \\ 4 & 4 & 4 \\ 6 & 1 & 0 \end{array} \right) \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} \quad \checkmark$

Wait, was that really any faster than elimination?

(1) Solve $\left(\begin{array}{ccc|c} 1 & 0 & 0 & 1 \\ 2 & 1 & 0 & 0 \\ 3 & -1 & 1 & 1 \end{array} \right)$ using substitution:

This elimination goes **down** instead of up, but it amounts to the same thing — just reorder the rows if you like. $\rightsquigarrow n^2$ flops

But the pivots are already = 1 (L is **unit** triangular)
so you don't have to do the n row scaling ops

$\rightsquigarrow n^2 - n$ flops

(2) Solve $\left(\begin{array}{ccc|c} 2 & 1 & 0 & 1 \\ 0 & 2 & 4 & -2 \\ 0 & 0 & 4 & -4 \end{array} \right)$ using substitution:

This is just Jordan substitution $\rightarrow n^2$ flops.

Total:

Solving $Ax=b$ using $A=LU$ takes
 $2n^2 - n \approx 2n^2$ flops

This turned elimination+substitution $\approx \frac{2}{3}n^3$ flops into $2n^2$.
Of course, you still have to compute $A=LU$:

Where does $A=LU$ come from?

If you can do Gaussian elimination without row swaps
then the only row operations you're allowed to do are

$$R_i \leftarrow cR_j, \quad i > j \quad \begin{array}{l} \text{(add a multiple of a row} \\ \text{to a row below it)} \end{array}$$

The corresponding elementary matrices are lower-unidular:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} R_3 \leftarrow 2R_3 \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix} = E$$

Let E_1, \dots, E_r be the elementary matrices for the row ops you performed in Gaussian elimination. They are lower-unidular.

$$A \xrightarrow[\text{ops}]{\text{row}} U = \text{REF}(A) \quad \text{means}$$

$$U = E_r E_{r-1} \cdots E_2 E_1 A$$

(left-multiplication by E_i does row operation i)

Since the E_i 's are lower-unital, so are

$$(E_r E_{r-1} \cdots E_2 E_1) \quad \text{and} \quad L = (E_r E_{r-1} \cdots E_2 E_1)^{-1} \\ = E_1^{-1} E_2^{-1} \cdots E_{r-1}^{-1} E_r^{-1}$$

$$\text{Then } LU = (E_r E_{r-1} \cdots E_2 E_1)^{-1} (E_r E_{r-1} \cdots E_2 E_1) A \\ = A \quad \checkmark$$

NB: $L = E_1^{-1} E_2^{-1} \cdots E_{r-1}^{-1} E_r^{-1}$ "keeps track" of the row operations you did in Gaussian elimination.

NB: $A = LU$ is a **matrix factorization**: it is a way of writing a matrix as a product of **simpler** matrices.
 → We'll learn many of these
 → They all make different calculations **faster**.

This also gives you a way of computing L & U .

$$L = E_1^{-1} E_2^{-1} \cdots E_{r-1}^{-1} E_r^{-1} = E_1^{-1} E_2^{-1} \cdots E_{r-1}^{-1} E_r^{-1} I_m$$

This means:

- start with I_m
- multiply by E_r^{-1} means **undo** the last row op.
- multiply by E_{r-1}^{-1} means **undo** the $(r-1)^{st}$ row op.
- ⋮
- multiply by E_2^{-1} means **undo** the 2nd row op.
- multiply by E_1^{-1} means **undo** the 1st row op.

This is L. (U is still the REF.)

Eg: $A = \begin{pmatrix} 2 & 1 & 0 \\ 4 & 4 & 4 \\ 6 & 1 & 0 \end{pmatrix}$

$R_2 \leftarrow 2R_1$ $\rightarrow \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 6 & 1 & 0 \end{pmatrix}$

$R_3 \leftarrow 3R_1$ $\rightarrow \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 0 & 2 & 0 \end{pmatrix}$

$R_3 \leftarrow R_2$ $\rightarrow \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 0 & 0 & 4 \end{pmatrix} = U$

Compute L:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \xrightarrow[R_3 \leftarrow R_2]{\text{undo last row op}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -1 & 1 \end{pmatrix} \xrightarrow[R_3 \leftarrow 3R_1]{\text{undo 2nd row op}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -3 & 1 \end{pmatrix}$$

$R_2 \leftarrow 2R_1$ $\rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & -3 & 1 \end{pmatrix} = L$



Here's a better way of doing the bookkeeping.

Computing $A=LU$ using the 2-Column Method:

(1) Start with a blank maxm "L" matrix on the left and your matrix A on the right.
→ Using the prescribed algorithm!

(2) Perform Gaussian elimination on A . For each row replacement $R_i \leftarrow cR_j$ put $-c$ in the (i,j) entry of the L matrix.

(3) Add 1's to the diagonal entries of L & 0's to the rest of the blank entries.

Now L is on the left and U on the right.

Eg:

	L	↔ 2 columns ↔	U
start blank →	$\begin{pmatrix} & \\ & \end{pmatrix}$		$\begin{pmatrix} 2 & 1 & 0 \\ 4 & 4 & 4 \\ 6 & 1 & 0 \end{pmatrix}$ ← start with A
$R_2 \leftarrow 2R_1$ ($c=2$)	$\begin{pmatrix} & \\ 2 & \end{pmatrix}$ ↑ (2,1) entry		$\begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 6 & 1 & 0 \end{pmatrix}$
$R_3 \leftarrow 3R_1$ ($c=3$)	$\begin{pmatrix} 2 & \\ 3 & \end{pmatrix}$ ↑ (3,1) entry		$\begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 0 & -2 & 0 \end{pmatrix}$
$R_3 \leftarrow R_2 + R_3$ ($c=1$)	$\begin{pmatrix} 2 & \\ 3 & \end{pmatrix}$ ↑ (3,2) entry -1 ↴		$\begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 4 \\ 0 & 0 & 4 \end{pmatrix} = U$

Fill in the blank entries of L:

$$\begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & -1 & 1 \end{pmatrix} \rightsquigarrow \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & -1 & 1 \end{pmatrix} = L$$

Computational Complexity

Finding $A = LU$ is just Gaussian elimination + some extra bookkeeping.

Computing $A = LU$:

$\approx \frac{2}{3}n^3$ flops

Solving $Ax = b$ given $A = LU$:

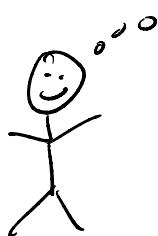
$\approx 2n^2$ flops

How does this help? Isn't this just as long as Gauss-Jordan elimination?

Yes, unless you have to solve $Ax = b$ for 1,000,000 different values of b! In this case, you just do elimination once and substitution 1,000,000 x.

Eg: If $n = 10^3$ and you have 10^6 b's:

- Gauss-Jordan 10^6 times is $10^6 \times \frac{2}{3} (10^3)^3 = \frac{2}{3} \cdot 10^{15}$ flops
- LU + substitution 10^6 times is $\frac{2}{3} (10^3)^3 + 10^6 \times 2 (10^3)^2 \approx 2 \cdot 10^{12}$ flops



If I want my computer to solve $Ax=b$ for a zillion values of b , I want to ask it for an LU decomposition first!

This is faster in SymPy too:

```
from sympy import *
from time import time

# This is the 10x10 matrix with 2's on the diagonal and 1's elsewhere
# eye(n) = nxn identity matrix; ones(n) = nxn matrix of 1's
# (multiply by 1.0 to force it to use floating point arithmetic)
A = (eye(10) + ones(10)) * 1.0
# This is the vector [1,1,1,1,1,1,1,1,1,1]
b = ones(10, 1) * 1.0

start = time()
# Compute LU decomposition
L, U, _ = A.LUdecomposition()
# Solve using substitution 1000 times
for _ in range(1000):
    U.upper_triangular_solve(L.lower_triangular_solve(b))
end = time()
print(end - start)
# "7.144780397415161" (seconds)

# Now solve using elimination 1000 times
start = time()
for _ in range(1000):
    A.solve(b)
end = time()
print(end - start)
# "48.048372983932495" (seconds)
# Roughly 10x slower!
```

What about inverses? Isn't it faster to solve for x by multiplying by A^{-1} ? $Ax=b \Leftrightarrow x=A^{-1}b$

- Computing A^{-1} uses $\approx \frac{4}{3}n^3$ flops.

- Multiplying $A^{-1}b$ uses $\approx 2n^2$ flops too!

- Computing A^{-1} ends up introducing more rounding errors.

Maximal Partial Pivoting

The system of equations

$$\begin{array}{l} x_2 = 1 \\ x_1 + x_2 = 2 \end{array} \quad \text{has one solution}$$

$$\begin{array}{l} x_1 = 1 \\ x_2 = 1 \end{array}$$

Let's tweak it just a little bit:

$$\begin{array}{l} -10^{-17}x_1 + x_2 = 1 \\ x_1 + x_2 = 2 \end{array}$$

Presumably this has one solution $x_1 \approx 1$, $x_2 \approx 1$.

$$\left(\begin{array}{cc|c} -10^{-17} & 1 & 1 \\ 1 & 1 & 2 \end{array} \right) \xrightarrow{R_2 + 10^{17}R_1} \left(\begin{array}{cc|c} -10^{-17} & 1 & 1 \\ 0 & 1+10^{17} & 2+10^{17} \end{array} \right)$$

$$\xrightarrow{R_2 \div 1+10^{17}} \left(\begin{array}{cc|c} -10^{-17} & 1 & 1 \\ 0 & 1 & \frac{2+10^{17}}{1+10^{17}} \end{array} \right) \xrightarrow{\frac{2+10^{17}}{1+10^{17}} = \frac{1+10^{17}}{1+10^{17}} + \frac{1}{1+10^{17}}} = 1 + \frac{1}{1+10^{17}}$$

$$\xrightarrow{R_1 \leftarrow R_2} \left(\begin{array}{cc|c} -10^{-17} & 0 & -\frac{1}{1+10^{17}} \\ 0 & 1 & \frac{1}{1+10^{17}} \end{array} \right) \xrightarrow{R_1 \times -10^{17}} \left(\begin{array}{cc|c} 1 & 0 & \frac{10^{17}}{1+10^{17}} \\ 0 & 1 & \frac{1}{1+10^{17}} \end{array} \right)$$

$$\text{So } x_1 = \frac{10^{17}}{1+10^{17}} \approx 1 \quad \text{and} \quad x_2 = 1 + \frac{1}{1+10^{17}} \approx 1 \quad \checkmark$$

Let's try this on the computer.

```

from sympy import *
# 1e-17 is 10^(-17)
A = Matrix([[[-1e-17, 1.0, 1.0],
             [1.0, 1.0, 2.0]]])

# This does R2 += 10^(17) R1
# (force sympy to use the smaller pivot)
A.row_op(1, lambda v, j: v + 1e17*A[0,j])
pprint(A)
# [-1e-17      1      1]
# [          0 1e17 1e17]

# Now do Jordan substitution
pprint(A.rref(pivots=False))
# [1 0 0]
# [0 1 1]
# This answer is numerically very wrong!

```

What went wrong?

Most computers represent decimal numbers in 64-bit floating point format.

https://en.wikipedia.org/wiki/IEEE_754

This amounts to ≈ 17 decimal digits of precision.

When the computer does

$$\begin{pmatrix} -10^{-17} & 1 & | & 1 \\ 1 & 1 & | & 2 \end{pmatrix} \xrightarrow{R_2 \leftarrow 10^{17}R_2} \begin{pmatrix} -10^{-17} & 1 & | & 1 \\ 0 & 1+10^{17} & | & 2+10^{17} \end{pmatrix}$$

It computes

$$1 + 10^{17} = \overbrace{1000000000000000000}^{17 \text{ digits}} 1$$

but it forgets the least significant (18th) digit:



$$\begin{aligned}
 2 + 10^{17} &= 100000000000000000002 \\
 &= 10000000000000000000 \\
 &\equiv 10^{17}
 \end{aligned}$$

Likewise:

So the computer does

$$\begin{pmatrix} -10^{-17} & 1 & | & 1 \\ 1 & 1 & | & 2 \end{pmatrix} \xrightarrow{R_2 + 10^{17}R_1} \begin{pmatrix} 10^{-17} & 1 & | & 1 \\ 0 & 10^{17} & | & 10^{17} \end{pmatrix}$$

$$\xrightarrow{R_1 \times 10^{17}} \begin{pmatrix} -10^{-17} & 1 & | & 1 \\ 0 & 1 & | & 1 \end{pmatrix} \xrightarrow{R_1 - R_2} \begin{pmatrix} -10^{-17} & 0 & | & 0 \\ 0 & 1 & | & 1 \end{pmatrix}$$

$$\xrightarrow{R_1 \times -10^{17}} \begin{pmatrix} 1 & 0 & | & 0 \\ 0 & 1 & | & 1 \end{pmatrix} \rightarrow \begin{matrix} x_1 = 0 \\ x_2 = 1 \end{matrix} \quad \times$$

Summary:

We had to divide R_1 by the 1st pivot $= -10^{-17}$
ie, we multiplied it by 10^{17}

then added it to R_2

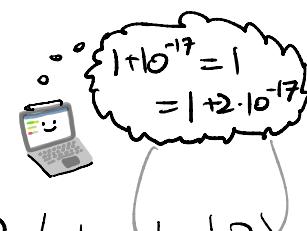
which had the effect of forgetting the rest of
the entries of R_2 .

How to fix this?

Row swap to choose the larger pivot!

$$\begin{pmatrix} -10^{-17} & 1 & | & 1 \\ 1 & 1 & | & 2 \end{pmatrix} \xrightarrow{R_1 \leftrightarrow R_2} \begin{pmatrix} 1 & 1 & | & 2 \\ -10^{-17} & 1 & | & 1 \end{pmatrix} \xrightarrow{R_2 + 10^{17}R_1} \begin{pmatrix} 1 & 1 & | & 2 \\ 0 & 1 & | & 1 \end{pmatrix}$$

$$\xrightarrow{R_1 - R_2} \begin{pmatrix} 1 & 0 & | & 1 \\ 0 & 1 & | & 1 \end{pmatrix} \rightarrow \begin{matrix} x_1 = 1 \\ x_2 = 1 \end{matrix}$$



✓ Much better.

There are several pivoting strategies for avoiding this kind of rounding error. Here is the simplest.

Gaussian Elimination with Maximal Partial Pivoting:

In step (a) of Gaussian elimination, perform a row swap so that (one of) the **largest** numbers in the column (in absolute value) becomes the new pivot.

NB: SymPy does this by default.

```
A = Matrix([[1e-17, 1.0, 1.0],  
           [1.0, 1.0, 2.0]])  
# Do Gauss-Jordan with MPP  
pprint(A.rref(pivots=False))  
# [1 0 1]  
# [0 1 1]
```

The story so far:

- Can solve $Ax=b$ faster with $A=LU$, which only works when you can do Gaussian elimination with **no row swaps**.
- Elimination is much more numerically accurate if you **row swap for every pivot**.

The best of both worlds is:

PALU Decompositions

Def: A **permutation matrix** is a product of elementary matrices for **row swaps**.

So if P is a permutation matrix then

$$\begin{aligned} PA &= (\text{do a bunch of row swaps on } A) \\ &= (\text{rearrange the rows of } A) \end{aligned}$$

PALU Decomposition:

Any matrix A has a factorization

$$PA = LU$$

where:

P is a permutation matrix

L is lower unit/ultral

U is an REF for A

Idea: First do all the row swaps on A (PA), then compute its LU decomposition ($PA = LU$).

Of course, you don't know which row swaps you'll need to do **before** doing elimination! Thankfully, this is taken care of with some slightly fancier bookkeeping.

Computing $PA = LU$ Using the 3-Column Method.

(1) Start with the $m \times m$ identity matrix "P" on the left, a blank $m \times m$ "L" matrix in the middle, and your matrix A on the right.

→ Using the prescribed algorithm!

(2) Do Gaussian Elimination on A.

- For each row replacement $R_i + cR_j$ put $-c$ in the (i,j) entry of the L matrix.
- For each row swap $R_i \leftrightarrow R_j$, swap the corresponding rows of L (including blank entries!) and P.

(3) Add 1's to the diagonal entries of L & 0's to the rest of the blank entries.

Now P is on the left, L in the middle, and U on the right.

Important: This only works if you do Gaussian elimination **as prescribed!** It will fail if you try to be more clever with your row operations.

Eg (PA=LU with Maximal Partial Pivoting):

Compute a PA=LU decomposition using MPP for

$$A = \begin{pmatrix} 1 & 1 & 1 \\ -10 & -20 & -30 \\ 5 & 15 & 10 \end{pmatrix}.$$

$$\begin{array}{c} P \\ \left(\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{array} \right) \\ \xrightarrow{\substack{R_1 \leftrightarrow R_2 \\ \text{choose largest pivot}}} \left(\begin{array}{ccc} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{array} \right) \end{array} \quad \begin{array}{c} L \\ \left(\begin{array}{ccc} & & \\ & & \\ & & \end{array} \right) \\ \xrightarrow{\substack{\text{---} \\ \text{---}}} \left(\begin{array}{ccc} \text{---} & \text{---} & \text{---} \\ \text{---} & \text{---} & \text{---} \end{array} \right) \end{array} \quad \begin{array}{c} U \\ \left(\begin{array}{ccc} 1 & 1 & 1 \\ -10 & -20 & -30 \\ 5 & 15 & 10 \end{array} \right) \\ \xrightarrow{\substack{\text{---} \\ \text{---}}} \left(\begin{array}{ccc} -10 & -20 & -30 \\ 1 & 15 & 10 \end{array} \right) \end{array}$$

$$\begin{array}{c} R_2 \leftarrow \frac{1}{10} R_1 \\ \xrightarrow{\substack{\text{---} \\ R_2 \leftarrow \frac{1}{2} R_1}} \left(\begin{array}{ccc} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{array} \right) \end{array} \quad \begin{array}{c} \left(\begin{array}{ccc} & & \\ -1/10 & & \\ -1/2 & & \end{array} \right) \\ \xrightarrow{\substack{\text{---} \\ \text{---}}} \left(\begin{array}{ccc} -1/10 & & \\ -1/2 & & \end{array} \right) \end{array} \quad \begin{array}{c} \left(\begin{array}{ccc} -10 & -20 & -30 \\ 0 & -1 & -2 \\ 0 & 5 & -5 \end{array} \right) \\ \xrightarrow{\substack{\text{---} \\ \text{---}}} \left(\begin{array}{ccc} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & -1 & -2 \end{array} \right) \end{array}$$

$$\begin{array}{c} R_2 \leftrightarrow R_3 \\ \xrightarrow{\substack{\text{choose largest pivot}}} \left(\begin{array}{ccc} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{array} \right) \end{array} \quad \begin{array}{c} \left(\begin{array}{ccc} & & \\ -1/2 & & \\ -1/10 & & \end{array} \right) \\ \xrightarrow{\substack{\text{---} \\ \text{---}}} \left(\begin{array}{ccc} -1/2 & & \\ -1/10 & & \end{array} \right) \end{array} \quad \begin{array}{c} \left(\begin{array}{ccc} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & -1 & -2 \end{array} \right) \\ \xrightarrow{\substack{\text{---} \\ \text{---}}} \left(\begin{array}{ccc} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & 0 & -3 \end{array} \right) \end{array}$$

$$\begin{array}{c} R_3 \leftarrow \frac{1}{5} R_2 \\ \xrightarrow{\substack{\text{---}}} \left(\begin{array}{ccc} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{array} \right) \end{array} \quad \begin{array}{c} \left(\begin{array}{ccc} & & \\ -1/2 & & \\ -1/10 & -1/5 & \end{array} \right) \end{array} \quad \begin{array}{c} \left(\begin{array}{ccc} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & 0 & -3 \end{array} \right) \end{array}$$

Fill in the blank entries of L:

$$\left(\begin{array}{ccc} & & \\ -1/2 & & \\ -1/10 & -1/5 & \end{array} \right) \rightsquigarrow \left(\begin{array}{ccc} 1 & 0 & 0 \\ -1/2 & 1 & 0 \\ -1/10 & -1/5 & 1 \end{array} \right) = L$$

Check:

$$\left(\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{array} \right) \left(\begin{array}{ccc} 1 & 1 & 1 \\ -10 & -20 & -30 \\ 5 & 15 & 10 \end{array} \right) = \left(\begin{array}{ccc} 1 & 0 & 0 \\ -1/2 & 1 & 0 \\ -1/10 & -1/5 & 1 \end{array} \right) \left(\begin{array}{ccc} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & 0 & -3 \end{array} \right)$$

Of course, this is only useful if we can use it to solve $Ax=b$:

How to solve $Ax=b$ using $PA=LU$:

- (0) Compute Pb . (rearrange the entries of b)
- (1) Solve $Ly = Pb$ using substitution.
- (2) Solve $Ux = y$ using substitution.

$$\text{Then } PAx = (LU)x = L(Ux) = Ly = Pb.$$

(Multiply both sides by P^{-1} $\Rightarrow Ax=b$.)

Eg: Solve $Ax=b$ where

$$A = \begin{pmatrix} 1 & 1 & 1 \\ -10 & -20 & -30 \\ 5 & 15 & 10 \end{pmatrix} \quad b = \begin{pmatrix} 0 \\ -10 \\ 10 \end{pmatrix}$$

using $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \\ -10 & -20 & -30 \\ 5 & 15 & 10 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ -1/2 & 1 & 0 \\ -1/10 & -1/5 & 1 \end{pmatrix} \begin{pmatrix} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & 0 & -3 \end{pmatrix}$

(0) $Pb = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ -10 \\ 10 \end{pmatrix} = \begin{pmatrix} -10 \\ 10 \\ 0 \end{pmatrix} \quad (\text{rearrange})$

(1) Solve $Ly = Pb$ using substitution:

$$\begin{pmatrix} 1 & 0 & 0 & | -10 \\ -1/2 & 1 & 0 & | 10 \\ -1/10 & -1/5 & 1 & | 0 \end{pmatrix} \xrightarrow{(\dots)} \begin{pmatrix} 1 & 0 & 0 & | -10 \\ 0 & 1 & 0 & | 5 \\ 0 & 0 & 1 & | 0 \end{pmatrix} \quad y = \begin{pmatrix} -10 \\ 5 \\ 0 \end{pmatrix}$$

(2) Solve $Ux=y$ using substitution:

$$\left(\begin{array}{ccc|c} -10 & -20 & -30 & -10 \\ 0 & 5 & -5 & 5 \\ 0 & 0 & -3 & 0 \end{array} \right) \xrightarrow{\text{...}} \left(\begin{array}{ccc|c} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{array} \right) \quad x = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}$$

Check:

$$\left(\begin{array}{ccc|c} 1 & 0 & 0 & -1 \\ -10 & -20 & -30 & 5 \\ 5 & 15 & 10 & 0 \end{array} \right) \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ -10 \\ 10 \end{pmatrix}$$



Computational Complexity

Finding $PA=LU$ is just Gaussian elimination + some extra bookkeepings, as with $A=LU$.

And solving $Ax=b$ using $PA=LU$ only had the extra step (0), which is just rearranging (no flops).

So the complexity is the same as $A=LU$.

Computing $PA=LU$: $\approx \frac{2}{3}n^3$ flops

Solving $Ax=b$ given $PA=LU$: $\approx 2n^2$ flops