Gaussian Elimination

This is how a computer solves systems of linear equations using elimination. Almost all questions in this class will reduce to this procedure! (The interesting part is how they do so.)

Def: Two matrices are row equivalent if you can get from one to the other using row operations.

NB: If augmented matrices are row equivalent then they have the same solution sets.

Algorithm (Gaussian Elimination/row reduction):

Input: Any matrix

Output: A row-equivalent matrix in REF.

Procedure:

(1a) If the first nonzero column has a zero entry at the top row swap so that the top entry is nonzero.

$$\begin{bmatrix} 0 & 4 & 3 & 3 \\ 1 & 1 & -1 & 3 \\ 5 & -3 & -6 & -6 \end{bmatrix} \xrightarrow{R_1 \leftrightarrow R_2} \begin{bmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 5 & -3 & -6 & -6 \end{bmatrix}$$

This is now the first pivot position.

(1b) Perform row replacements to clear all entries below the first pivot.

$$\begin{bmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 5 & -3 & -6 & -6 \end{bmatrix} \xrightarrow{R_3 -= 5R_1} \begin{bmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 0 & -8 & -1 & -21 \end{bmatrix}$$

Now ignore the row & column with the first pivot and recurse into the submatrix below and to the right:

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

(2a) If the first nonzero column has a zero entry at the top row swap so that the top entry is nonzero.

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

second pivot

(Not applicable to this matrix)

(2b) Perform row replacements to clear all entries below the second pivot.

$$\begin{bmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 0 & -8 & -1 & -21 \end{bmatrix} \xrightarrow{R_3 += 2R_2} \begin{bmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 0 & 0 & 5 & -15 \end{bmatrix}$$

Doesn't mess up the 1st column!

etc. (recurse)

In our example, the recursion has terminated:

$$\begin{bmatrix} 1 & 1 & -1 & 3 \\ 0 & 4 & 3 & 3 \\ 0 & 0 & 5 & -15 \end{bmatrix} \quad \text{is in REF!}$$

Important: If you want to apply this algorithm to an augmented matrix, just delete the augmentation line (pretend it's not augmented).

Demo: Gauss-Jordan slideshow

Use Rabinoff's Reliable Row Reducer on the HW!
↳ use the Sage cell when this gets easy

Jordan Substitution

This is the back-substitution procedure.
It is necessary when you have ∞ solutions.
It puts a matrix into the following form:

Def: A matrix is in reduced row echelon form (RREF) if:

(1–2) It is in REF
(3) All pivots are equal to 1.
(4) A pivot is the only nonzero entry in its column.

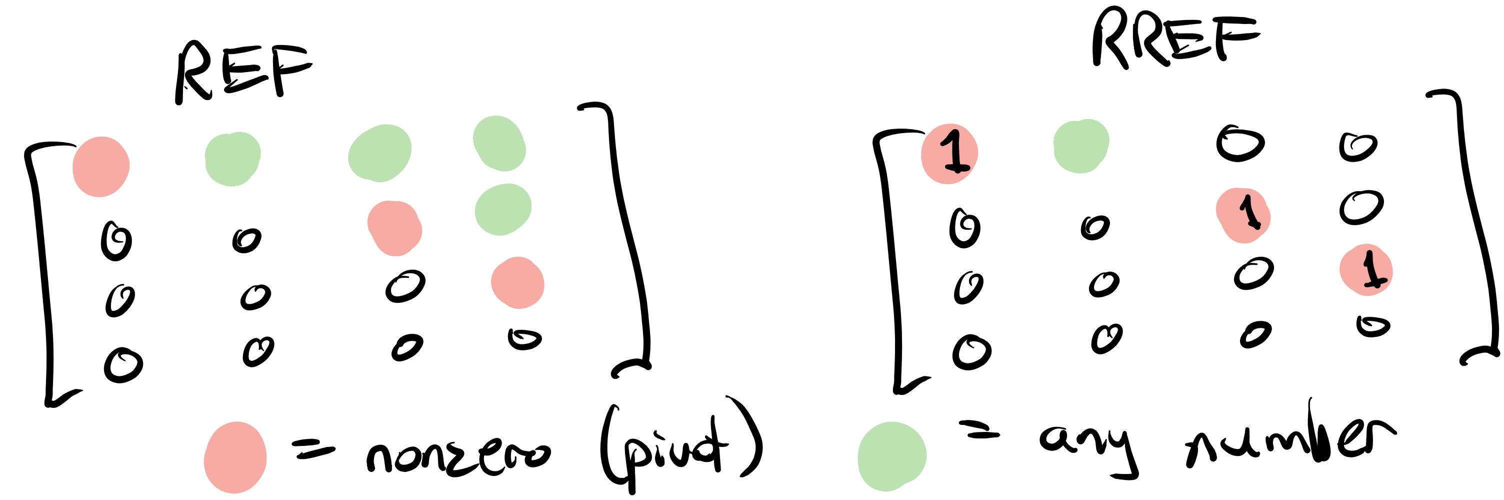

Eg: $\left[\begin{array}{ccc|c} 1 & 2 & 3 & 6 \\ 0 & -5 & -10 & -20 \\ 0 & 0 & 10 & 30 \end{array}\right]$ is in REF. How to put into RREF?

Do back substitution!

Row Operations

$$\left[\begin{array}{ccc|c} 1 & 2 & 3 & 6 \\ 0 & -5 & -10 & -20 \\ 0 & 0 & 10 & 30 \end{array}\right]$$

(scale so this is 1)

$R_3 \mathrel{\div}= 10$ { solve for x_3

$$\left[\begin{array}{cccc} 1 & 2 & 3 & 6 \\ 0 & -5 & -10 & -20 \\ 0 & 0 & 1 & 3 \end{array}\right]$$

(kill these)

$R_1 \mathrel{-}= 3R_3$

$R_2 \mathrel{+}= 10R_3$

{ substitute $x_3 = 3$ into R_1 & R_2 then move the constants to the RHS

Back-Substitution

$$x_1 + 2x_2 + 3x_3 = 6$$
$$-5x_2 - 10x_3 = -20$$
$$10x_3 = 30$$

$$x_1 + 2x_2 + 3x_3 = 6$$
$$-5x_2 - 10x_3 = -20$$
$$x_3 = 3$$

$$\begin{bmatrix} 1 & 2 & 0 & -3 \\ 0 & -5 & 0 & 10 \\ 0 & 0 & 1 & 3 \end{bmatrix} \qquad \begin{aligned} x_1 + 2x_2 &= -3 \\ -5x_2 &= 10 \\ x_3 &= 3 \end{aligned}$$

(scale so this is 1)

$R_2 \div= -5$ — solve for x_2

$$\begin{bmatrix} 1 & 2 & 0 & -3 \\ 0 & 1 & 0 & -2 \\ 0 & 0 & 1 & 3 \end{bmatrix} \qquad \begin{aligned} x_1 + 2x_2 &= -3 \\ x_2 &= -2 \\ x_3 &= 3 \end{aligned}$$

(kill this)

$R_1 -= 2R_2$ — substitute $x_2 = -2$ into R_1 then move the constants to the RHS

$$\begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & -2 \\ 0 & 0 & 1 & 3 \end{bmatrix} \qquad \begin{aligned} x_1 &= 1 \\ x_2 &= -2 \\ x_3 &= 3 \end{aligned}$$

This is in RREF:

$$\left[\begin{array}{ccc|c} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & -2 \\ 0 & 0 & 1 & 3 \end{array}\right] \rightsquigarrow \begin{aligned} x_1 &= 1 \\ x_2 &= -2 \\ x_3 &= 3 \end{aligned}$$

Solved ✓

Upshot: Jordan substitution is exactly back-substitution.

Demo: Gauss-Jordan slideshow, cont'd

Algorithm (Jordan Substitution):

Input: A matrix in REF

Output: The row-equivalent matrix in RREF.

Procedure:

Loop, starting at the last pivot:

(a) Scale the pivot row so the pivot = 1.

(b) Use row replacements to kill the entries above that pivot.

"theorem"

Thm: The RREF of a matrix is unique.

In other words, if you start with a matrix, do any legal row operations at all, and end with a matrix in RREF, then it's the same matrix that Gauss-Jordan will produce.

↳ Gaussian elimination + Jordan substitution.

NB: Jordan substitution gives you a RREF matrix with the same pivots. So uniqueness of RREF implies uniqueness of pivot positions.

Computational Complexity

How much computer time does Gauss-Jordan take?

Gaussian Elimination on an $n \times n$ matrix takes:

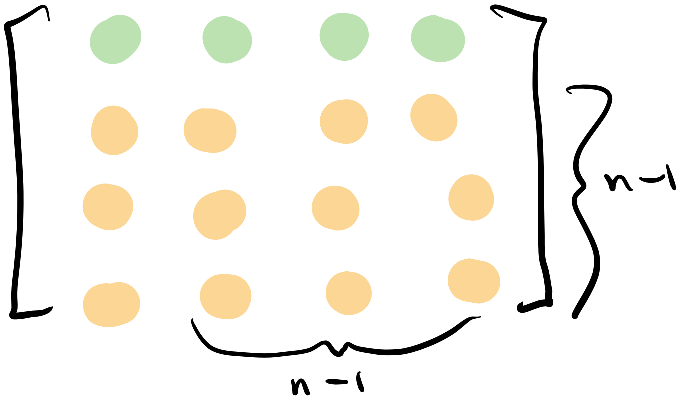

Step 1: Each row replacement requires $n-1$ multiplications & $n-1$ additions. (no computation in 1st col: just write "0")

Do for $n-1$ lower rows:

$$\begin{array}{ll} (n-1)(n-1) & \text{mult} \\ (n-1)(n-1) & \text{add} \\ \hline 2(n-1)^2 & \end{array}$$

flops = floating point operations

Step 2: Each row replacement requires $n-2$ multiplications & $n-2$ additions. Must do this for $n-2$ remaining rows

$$\begin{array}{ll} (n-2)(n-2) & \text{mult} \\ +\ (n-2)(n-2) & \text{add} \\ \hline 2(n-2)^2 & \text{flops} \end{array}$$

etc.

Total: $2\,[\underbrace{(n-1)^2 + (n-2)^2 + \cdots + 1^2}_{\text{pyramidal number}}]$

$$= 2 \cdot \frac{n(n-1)(2n-1)}{6} \approx \frac{2}{3}n^3 \text{ flops}$$

Back-Substitution

● x_n = ● — 1 mult = 1 flop

● x_{n-1} + ● x_n = ● — 2 mult, 1 add = 3 flops

(substitute x_n, × ●, subtract, ÷ ●)

● x_{n-2} + ● x_{n-1} + ● x_n = ● — 3 mult, 2 add = 5 flops

(substitute x_n & x_{n-1}, × ●, × ●, subtract, ÷ ●)

● x_1 + ⋯ + ● x_n = ● — n mult, $(n-1)$ add = $2n-1$ flops

Total: $1+3+5+\cdots+(2n-1) = n^2$ flops

NB: $\frac{2}{3}n^3$ is a lot more than n^2!

For a 1000×1000 matrix, $\frac{2}{3}n^3 \approx \frac{2}{3}$ gigaflops but $n^2 = 1$ megaflop. If we want to solve $Ax=b$ for 1000 values of b, doing elimination each time takes $\frac{2}{3}$ teraflops!

Inverse Matrices

Question: When solving $Ax=b$, when can we "divide by A"?

If "$x=\frac{b}{A}$" makes sence, then $Ax=b$ has exactly one solution "$x=\frac{b}{A}$" for every b.

This means RREF(A|b) looks like this:

$$\left(\begin{array}{ccc|c} 1 & 0 & 0 & b_1 \\ 0 & 1 & 0 & b_2 \\ 0 & 0 & 1 & b_3 \end{array}\right)$$

Def: An $n\times n$ (square!) matrix A is invertible if there exists another $n\times n$ matrix B such that $AB = I_n = BA$.

$I_n = \begin{bmatrix} 1 & & 0 \\ & \ddots & \\ 0 & & 1 \end{bmatrix}$ $n\times n$ identity matrix

Otherwise it's called singular.

Notn: $B=A^{-1}$, called the inverse of A.

Eg: $A=\begin{bmatrix} 2 & 3 \\ 1 & 2 \end{bmatrix}$ $B=\begin{bmatrix} 2 & -3 \\ -1 & 2 \end{bmatrix}$

$AB=\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ $BA=\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$

$\Rightarrow B=A^{-1}$

Eg: $A=\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}\underbrace{\begin{bmatrix} a & b \\ c & d \end{bmatrix}}_{\text{"B"}} = \begin{bmatrix} a & b \\ 0 & 0 \end{bmatrix} \neq I_2$

so A is singular (non-invertible).

Remark: Since $AB \neq BA$ in general, you have to require $AB = I_n = BA$ a priori. But:

Fact: If A and B are $n \times n$ matrices and $AB = I_n$ or $BA = I_n$, then $B = A^{-1}$.

So the definition above is a bit pedantic...

Remark: A non-square matrix does not admit both a left- and right-inverse, so not invertible. (Can't solve $AB = I_m$ and $CA = I_n$ unless A is square.)

This is why we only treat invertibility of square matrices.

Fact: $(A^{-1})^{-1} = A$

because $AB = I_n$ means $B = A^{-1}$ and $A = B^{-1}$

Fact: If A & B are invertible, then so is AB, and $(AB)^{-1} = \cancel{A^{-1}B^{-1}}\ B^{-1}A^{-1}$.

Check: $(AB)(B^{-1}A^{-1}) = A(BB^{-1})A^{-1} = A(I_n)A^{-1} = AA^{-1} = I_n$ ✓

NB: Why not $(AB)^{-1} = A^{-1}B^{-1}$? Let's check:

$$(AB)(A^{-1}B^{-1}) = ABA^{-1}B^{-1} = ???$$

(no cancellation — can't re-order the terms!)

Thm: Let A be an $n \times n$ matrix. The Following Are Equivalent: (TFAE) ← either all are true or all are false

(1) A is invertible

(2) The RREF of A is I_n ← coefficient matrix!

(3) A has a pivot in every row/every column. (A has n pivots)

We'll see why a bit later.

Eg: $A = \begin{bmatrix} 2 & 3 \\ 1 & 2 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = I_2$ invertible

● = pivots

Eg: $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ is in RREF, $\neq I_2$ singular

How do you compute the inverse?

Algorithm (Matrix Inversion): ← why does this work? Next time.

Input: A square matrix.

Output: The inverse matrix, or "singular"

Procedure:

(a) Form the augmented matrix $[A \mid I_n]$

(b) Run Gauss-Jordan on $[A \mid I_n]$.

(c) If the output is $[I_n \mid B]$ then $B = A^{-1}$.
Otherwise A is singular.

Eg: Compute $\begin{bmatrix} 2 & 3 \\ 1 & 2 \end{bmatrix}^{-1}$.

$$\left[\begin{array}{cc|cc} 2 & 3 & 1 & 0 \\ 1 & 2 & 0 & 1 \end{array}\right] \xrightarrow{R_2 -= \frac{1}{2}R_1} \left[\begin{array}{cc|cc} 2 & 3 & 1 & 0 \\ 0 & 1/2 & -1/2 & 1 \end{array}\right]$$

$$\xrightarrow{R_2 \times= 2} \left[\begin{array}{cc|cc} 2 & 3 & 1 & 0 \\ 0 & 1 & -1 & 2 \end{array}\right]$$

$$\xrightarrow{R_1 -= 3R_2} \left[\begin{array}{cc|cc} 2 & 0 & 4 & -6 \\ 0 & 1 & -1 & 2 \end{array}\right]$$

$$\xrightarrow{R_1 \div= 2} \left[\begin{array}{cc|cc} 1 & 0 & 2 & -3 \\ 0 & 1 & -1 & 2 \end{array}\right]$$

This has the form $\left[\begin{array}{c|cc} I_2 & 2 & -3 \\ & -1 & 2 \end{array}\right]$

So $\begin{bmatrix} 2 & 3 \\ 1 & 2 \end{bmatrix}^{-1} = \begin{bmatrix} 2 & -3 \\ -1 & 2 \end{bmatrix}$.

Eg: Compute $\begin{bmatrix} 1 & 3 \\ -1 & -3 \end{bmatrix}^{-1}$.

$$\left[\begin{array}{cc|cc} 1 & 3 & 1 & 0 \\ -1 & -3 & 0 & 1 \end{array}\right] \xrightarrow{R_2 += R_1} \left[\begin{array}{cc|cc} 1 & 3 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{array}\right]$$

$$\xrightarrow{R_1 -= R_2} \left[\begin{array}{cc|cc} 1 & 3 & 0 & -1 \\ 0 & 0 & 1 & 1 \end{array}\right]$$

But this is in RREF and it does not have the form $[I_2 \mid B] \Rightarrow$ singular.

NB: We knew this after the first step: no pivot in the second column.

Actually there's a shortcut for 2x2 matrices:

Fact: $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is invertible $\iff ad-bc \neq 0$,
in which case

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad-bc}\begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$$

Eg: $$\begin{bmatrix} 2 & 3 \\ 1 & 2 \end{bmatrix}^{-1} = \frac{1}{2\cdot 2-3\cdot 1}\begin{bmatrix} 2 & -3 \\ -1 & 2 \end{bmatrix}$$

Check: $$\frac{1}{ad-bc}\begin{bmatrix} d & -b \\ -c & a \end{bmatrix}\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$
$$= \frac{1}{ad-bc}\begin{bmatrix} ad-bc & bd-bd \\ ac-ac & ad-bc \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \checkmark$$

What is this good for?
Suppose A is invertible. Let's solve $Ax=b$.

$$Ax=b \iff A^{-1}(Ax) = A^{-1}b$$
$$\iff (A^{-1}A)x = A^{-1}b$$
$$\iff I_n x = A^{-1}b \iff x = A^{-1}b$$

For invertible A:

$$Ax=b \iff x=A^{-1}b$$

In particular, $Ax=b$ has exactly one solution for any b, and we have an expression for b in terms of x

Eg: Solve
$$2x_1+3x_2=b_1$$
$$x_1+2x_2=b_2.$$

$$A=\begin{bmatrix}2&3\\1&2\end{bmatrix} \qquad A^{-1}=\begin{bmatrix}2&-3\\-1&2\end{bmatrix}$$

$$A\begin{bmatrix}x_1\\x_2\end{bmatrix}=\begin{bmatrix}b_1\\b_2\end{bmatrix} \iff \begin{bmatrix}x_1\\x_2\end{bmatrix}=A^{-1}\begin{bmatrix}b_1\\b_2\end{bmatrix}=\begin{bmatrix}2&-3\\-1&2\end{bmatrix}\begin{bmatrix}b_1\\b_2\end{bmatrix}$$

$$\iff \begin{aligned} x_1 &= 2b_1-3b_2 \\ x_2 &= -b_1+2b_2 \end{aligned}$$

So if you want to solve
$$2x_1+3x_2=3$$
$$x_1+2x_2=4$$

$$\Rightarrow x_1=2(3)-3(4)=-6$$
$$x_2=-(3)+2(4)=5$$