Applications

- Robot path planning
- Network flows (transportation, circuits, communication)
- Equilibria (physics, economics)

Topics

- Pivoting and permutations (ALA 1.4)
  - singular & nonsingular matrices, uniqueness of sol'ns
  - permutation matrices
  - permuted LU factorization

- Matrix inverses (ALA 1.5)
  - Gauss-Jordan Elimination to compute A^{-1}
  - $\underline{x} = A^{-1}\underline{b}$

- General linear systems (ALA 1.8)
  - row echelon form
  - the rank of a matrix
  - basic and free variables
  - Conditions for no, unique, or infinite solutions to $A\underline{x} = \underline{b}$.

Pivoting & Permutations

What if A is not regular? What if a zero appears at the current pivot position?

Example:

$$\begin{aligned} 2y + z &= 2 \\ 2x + 6y + z &= 7 \\ x + y + 4z &= 3 \end{aligned} \longrightarrow M = \left[\begin{array}{ccc|c} 0 & 2 & 1 & 2 \\ 2 & 6 & 1 & 7 \\ 1 & 1 & 4 & 3 \end{array}\right]$$

The zero in the $(1,1)$ position means we don't have a first pivot.

"Problem" is because x does not appear in 1^{st} equation.

Actually a bonus! Only y & z appear in 1^{st} equation. Let us interchange the first two rows of M (equivalent to swapping 1^{st} & 2^{nd} equations):

$$M_1 = \left[\begin{array}{ccc|c} 2 & 6 & 1 & 2 \\ 0 & 2 & 1 & 7 \\ 1 & 1 & 4 & 3 \end{array}\right]$$

Clearly this does not change the solution to the equations, and is allowed!

Tool: Interchange two rows of the matrix. (T2)

Now we can proceed as before, to ultimately recover the triangular system augmented matrix:

$$N = \left[\begin{array}{ccc|c} 2 & 6 & 1 & 7 \\ 0 & 2 & 1 & 2 \\ 0 & 0 & 9/2 & 3/2 \end{array}\right]$$

$$\underbrace{\qquad\qquad}_{U} \qquad \underbrace{\quad}_{\underline{c}}$$

The pivots are $2, 2, 9/2$ and solving $U\underline{x} = \underline{c}$ via Back Substitution yields a solution $x = \frac{5}{6}, y = \frac{5}{6}, z = \frac{1}{3}$.

We need to use row interchanges (T2) whenever a zero shows up in a diagonal pivot position.

When a matrix can be reduced to upper triangular form by Gaussian Elimination with pivoting, we call it nonsingular.

Theorem: A linear system $A\underline{x} = \underline{b}$ has a unique solution for every choice of RHS $\underline{b}$ if and only if the coefficient matrix A is square and nonsingular.

Our discussion thus far shows the "if" direction: if Gaussian Elimination with pivoting succeeds, the solution exists and is unique. We will return to the "only if" direction later.

NUMERICAL NOTE

Sometimes we want to interchange rows when the pivot is very small but nonzero. Here is a simple example:

$$\begin{aligned} 0.01x + 1.6y &= 32.1 \\ x + 0.6y &= 22 \end{aligned}$$

We can solve this to get the exact solution $x = 10$ & $y = 20$

Now suppose we have a calculator that only retains 3 digits of accuracy (modern computers of course have much higher, but still finite, accuracy. We can modify this example to produce similar issues there too).

Let us use our calculator to solve our system using the augmented matrix:

$$\left[\begin{array}{cc|c} .01 & 1.6 & 32.1 \\ 1 & .6 & 22 \end{array}\right]$$

Using the (1,1) as our pivot, we subtract $100\cdot(\text{row }1)$ from (row 2):

$$\left[\begin{array}{cc|c} .01 & 1.6 & 32.1 \\ 0 & -159.4 & -3188 \end{array}\right]$$

But our calculator only keeps 3 digits around so we actually get

$$\left[\begin{array}{cc|c} .01 & 1.6 & 32.1 \\ 0 & -159.0 & -3190 \end{array}\right]$$

The solution by Back substitution is $y = 3190/159 \approx 20.1$
$x = 100(32.1 - 32.16) \approx -10$.

A small error in y (only 0.1 off!) completely messes up our solution for x!

The culprit is our very small pivot .01: you can check that if we first swapped the equations & then solved the system, we would get approximate solutions of $y \simeq 20.1$ & $x \simeq 9.9$.

Permutations and Permutation Matrices

Just as with adding rows together, we want to encode swapping rows via matrix multiplication. Following our earlier template for constructing matrices associated with row operations, we define a permutation matrix as one obtained by interchanging rows of the identity matrix.

For example, if I want to swap rows 1 & 2 of a matrix with three rows I can left multiply by

$$P = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \Rightarrow \underbrace{\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{P} \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix} = \begin{bmatrix} 4 & 5 & 6 \\ 1 & 2 & 3 \\ 7 & 8 & 9 \end{bmatrix}.$$

An alternative definition for a permutation matrix is a square matrix where each column has all 0 entries except a single 1 entry, and each row has all 0 entries except a single 1 entry.

Some key properties of permutation matrices:

- If P_1 & P_2 are permutation matrices, then so are the products $P_1 P_2$ & $P_2 P_1$.
- In general, $P_1 P_2 \neq P_2 P_1$ (the order of swapping rows matters.)
- If P_1 is a permutation matrix that only swaps one row with another, then $P_1 P_1 = I$. This is NOT TRUE for general permutation matrices though! (why?)

The Permuted LU-factorization

We know every nonsingular A can be reduced to upper triangular form by adding & interchanging rows. If we do all of the row swaps ahead of time, say using a permutation matrix P, then we end up with a regular matrix PA that admits an LU-factorization. We can define a permuted LU-factorization:

$$PA = LU$$

swap the rows of A to make PA regular

apply the LU-factorization to the regular matrix PA.

Applying these ideas is not hard, but also not particularly interesting. See Example 1.12 in ALA 1.4 & online notes for some worked examples.

Using the permuted LU-factorization along with Forward & Back Substitution, we can solve $A\underline{x} = \underline{b}$ whenever A is nonsingular:

① First apply permutation to both sides: $PA\underline{x} = P\underline{b} = \hat{\underline{b}}$ ($\hat{\underline{b}}$ is just a new name for $P\underline{b}$)

② Now solve the two triangular systems:

$$L\underline{z} = \hat{\underline{b}} \quad \text{and} \quad U\underline{x} = \underline{c}$$

by Forward & Back Substitution, as before.

Matrix Inverses

Suppose a, x, and b are scalars, and we want to solve $ax=b$. If $a \neq 0$, then we can multiply both sides by $a^{-1} = \frac{1}{a}$ to obtain $x = a^{-1}b = \frac{b}{a}$.

We will define the matrix analogue of $a^{-1} = 1/a$, called the matrix inverse. As was the case with matrix arithmetic, some aspects are pretty similar to normal division, & others are a bit different.

For a square $n \times n$ matrix A, an $n \times n$ matrix X is called the inverse of A if

$$XA = AX = I_n.$$

If it exists, we denote the inverse of A with A^{-1}.

WARNING: Not every square matrix has an inverse, just like not every scalar does: $0^{-1} = 1/0$ is not defined since $0x = 1$ has no solution.

WARNING: In general, there is no obvious way to guess the entries of A^{-1} by looking at the entries of A.

The inverse of a matrix is typically more useful in theory than in practice: why this is the case will become clear as the semester progresses. As such though, we will not spend too much time on computing inverses — it is very rare you will ever need to outside of a linear algebra class!

Example: 2×2 matrices

Let us come up with a general formula for the inverse of a 2×2 matrix.

Let $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ & $X = \begin{bmatrix} x & y \\ z & w \end{bmatrix}$. For X to be the inverse of A it must satisfy $AX = I$

$$AX = \begin{bmatrix} a & b \\ c & d \end{bmatrix}\begin{bmatrix} x & y \\ z & w \end{bmatrix} = \begin{bmatrix} ax+bz & ay+bw \\ cx+dz & cy+dw \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = I$$

This holds if and only if x, y, z, w satisfy the linear system

$$\begin{aligned} ax + bz &= 1 \qquad & ay + bw &= 0 \\ cx + dz &= 0 \qquad & cy + dw &= 1 \end{aligned}$$

Solving by Gaussian Elimination, we find:

$$x = \frac{d}{ad-bc}, \quad y = \frac{-b}{ad-bc}, \quad z = \frac{-c}{ad-bc}, \quad w = \frac{a}{ad-bc}$$

provided that the common denominator $ad-bc \neq 0$. In that case, the matrix

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

is a candidate for the inverse of A. We also check (by matrix mult.) that $XA = I$ also holds.

Therefore $X = A^{-1}$ is the inverse of A.

Example Elementary Matrix for Row Addition

We saw previously that the inverse elementary matrix for row addition can be easily computed. For example adding $2\cdot$(row 1) to (row 3) is encoded via

$$E = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{bmatrix}$$

and undoing it by subtracting $2\cdot$(row 1) from (row 3) is encoded by

$$L = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -2 & 0 & 1 \end{bmatrix}.$$

You can check that $L = E^{-1}$, i.e., that $LE = EL = I$.

Idea: the inverse of a matrix A "undoes" the changes A does to a vector via matrix-vector multiplication Ax.

$$\boxed{\underline{x} \xrightarrow{A} A\underline{x} \xrightarrow{A^{-1}} A^{-1}A\underline{x} = \underline{x}}$$

Example Single Row Swap.

The matrix $P = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ swaps (row 1) & (row 2) for a three row matrix.

P is its own inverse! $PP = P^2 = I$. The idea above gives an easy intuition as to why this is true (but you can also check directly by tediously computing the matrix product).

IMPORTANT FACT: A square matrix has an inverse if and only if it is non singular.

We will establish this rigorously later, but let's revisit our earlier scalar analogy. The reciprocal a^{-1} exists if and only if the scalar equation $ax = b$ has a unique solution $x = a^{-1}b$. This is true if and only if $a \neq 0$.

Similarly, the inverse A^{-1} exists if and only if the linear system $A\underline{x} = \underline{b}$ has a unique solution $\underline{x} = A^{-1}\underline{b}$.

<u>Some properties of matrix inverses</u>

- The inverse of a square matrix, if it exists, is unique:

  Suppose both X & Y satisfy
  $$AX = XA = I$$
  $$AY = YA = I$$

  Then $X = X \cdot I = X(AY) = (XA)Y = I \cdot Y = Y$.

- If A is invertible, so is A^{-1}, and $(A^{-1})^{-1} = A$:

  If A^{-1} is the inverse of A, then $AA^{-1} = A^{-1}A = I$. But these are the same conditions for A to be the inverse of A^{-1}!

  Using our idea above, this says that the way to "undo" (invert) "undoing A" ($x \to A^{-1}x$) is to apply A ($A^{-1}x \to A(A^{-1})x = x$)

  $$x \xrightarrow{A^{-1}} A^{-1}x \xrightarrow{A} AA^{-1}x.$$

- If A & B are invertible matrices of the same size, then their product AB is invertible, and
  $$(AB)^{-1} = B^{-1}A^{-1}$$

  NOTE THAT THE ORDER IS REVERSED IN THE INVERSE!!!

  Let $X = B^{-1}A^{-1}$, then $X(AB) = B^{-1}\underbrace{A^{-1}A}_{=I}B = \underbrace{B^{-1}B}_{=I} = I$.

  and $(AB)X = A\underbrace{BB^{-1}}_{=I}A^{-1} = \underbrace{AA^{-1}}_{=I} = I$.

  If we use our idea, this makes sense, we need to undo transformations in the right order:

  $$x \underset{\text{apply } B}{\xrightarrow{B}} Bx \underset{\text{apply } A}{\xrightarrow{A}} ABx \underset{\text{undo } A}{\xrightarrow{A^{-1}}} A^{-1}ABx = Bx \underset{\text{undo } B}{\xrightarrow{B^{-1}}} B^{-1}Bx = x.$$

Gauss-Jordan Elimination

Gauss-Jordan Elimination (GJE) is the principal algorithm for computing inverses on nonsingular matrices.

Key fact: for a square matrix A, we only need to solve $AX = I$ (since $AX = XA = I$, for some matrix $X = A^{-1}$)

** NOTE: for some matrices, we can only find a right inverse X_R satisfying $AX_R = I$; similarly, some matrices only admit a left inverse $X_L A = I$. This is why we emphasize that an inverse X must satisfy both $XA = I$ & $AX = I$, even if checking one condition is enough. **

We already have the tools needed to solve $AX = I$! First, we define the $n \times 1$ unit vectors $\underline{e_i}$:

$$\underline{e_1} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \underline{e_2} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \ldots, \quad \underline{e_n} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}$$

as the vectors with one entry of 1 in i^{th} position, and zeros elsewhere. These also happen to be the columns of the $n \times n$ identity matrix I_n:

$$I_n = [\underline{e_1} \ \underline{e_2} \ \cdots \ \underline{e_n}].$$

Letting the columns of X be $\underline{x_1}, \underline{x_2}, \ldots, \underline{x_n}$, so that $X = [\underline{x_1} \ \underline{x_2} \cdots \underline{x_n}]$, we have

$$AX = I \iff A[\underline{x_1} \ \underline{x_2} \cdots \underline{x_n}] = [A\underline{x_1} \ A\underline{x_2} \ \cdots \ A\underline{x_n}]$$
$$= [\underline{e_1} \ \underline{e_2} \ \cdots \ \underline{e_n}]$$

i.e., the matrix equation $AX = I$ is equivalent to n linear systems

$$A\underline{x_1} = \underline{e_1}, \quad A\underline{x_2} = \underline{e_2}, \quad \ldots, \quad A\underline{x_n} = \underline{e_n}.$$

A key feature here is that all n linear systems have the same coefficient matrix A! We should take advantage of that!

With that in mind, instead of forming n augmented matrices $M_1 = [A \mid \underline{e_1}]$, $M_2 = [A \mid \underline{e_2}]$, $\ldots$, $M_n = [A \mid \underline{e_n}]$, let's just form a single large one:

$$M = [A \mid \underline{e}_1\ \underline{e}_2 \cdots \underline{e}_n] = [A \mid I].$$

We can then simultaneously apply our row operations (swaps & adding) to reduce A to upper triangular form:

$$M = [A \mid I] \longrightarrow N = [U \mid C]$$

which is equivalent to reducing the original n linear systems to $U\underline{x}_1 = \underline{c}_1$, $U\underline{x}_2 = \underline{c}_2, \ldots, U\underline{x}_n = \underline{c}_n$. We could then solve for the $\underline{x}_i$ via Back substitution.

A more common version of GJE continues to apply row operations to fully reduce the augmented matrix until it is of the form $[I \mid X]$, where here $X = A^{-1}$ is the desired matrix inverse.

Both U and I have zeros below the diagonal, but I has 1's along its diagonal; U has nonzero pivots. We need another tool!

Tool: we can multiply a row by a nonzero scalar.

This is equivalent to scaling an equation by a nonzero number; clearly this does not change the solution.

Q: Can you encode this row operation via a left multiplication by an elementary matrix?

For a worked example, see the online notes. Remember, this is no different than solving a linear system, we just now keep track of n RHS vectors instead of one.

General Linear Systems

So far we have treated only "nice" linear systems defined by nonsingular coefficient matrices; these are "nice" because they have a unique solution.

We now turn to systems w/ m equations in n unknowns, including non-square $(m \neq n)$ & singular coefficient matrices.

The good news is that we already have the tools we need to solve the general case! But there are some minor changes, which we will go over now.

Our starting point is a generalization of upper triangular matrices in "staircase" form. A $m \times n$ matrix U is in row echelon form if it has the following form:

$$U = \begin{pmatrix} \circledast & * & \cdots & * & * & \cdots & * & * & \cdots & \cdots & * & * & * & \cdots & * \\ 0 & 0 & \cdots & 0 & \circledast & \cdots & * & * & \cdots & \cdots & * & * & * & \cdots & * \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & \circledast & \cdots & \cdots & * & * & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & & \ddots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & \cdots & \cdots & 0 & \circledast & * & \cdots & * \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & \cdots & \cdots & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & \cdots & \cdots & 0 & 0 & 0 & \cdots & 0 \end{pmatrix} \begin{array}{l} \left.\begin{matrix} \\ \\ \\ \\ \\ \end{matrix}\right\} r \text{ rows} \\ \left.\begin{matrix} \\ \\ \\ \end{matrix}\right\} m-r \text{ rows} \end{array}$$

$\circledast$ are the pivots, and must be nonzero

$*$ are arbitrary, possibly zero, entries.

The first r rows each contain a single pivot, but not all columns need have one.

All entries below the "staircase" are zero

The last $m-r$ rows are all zeros, and have no pivots.

Here is an example of a matrix in row echelon form:

$$\begin{bmatrix} 3 & 1 & 0 & 4 & 5 & -7 \\ 0 & -1 & -2 & 1 & 8 & 0 \\ 0 & 0 & 0 & 0 & 2 & -4 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{array}{l} \left.\begin{matrix} \\ \\ \\ \end{matrix}\right\} r=3 \\ \left.\begin{matrix} \\ \end{matrix}\right\} m-r=1 \end{array}$$

The pivots are circled: they are 3, -1, and 2.

EXCEPTION: in rare cases, the matrix may start with all zero columns for example

$$\begin{bmatrix} 0 & 0 & 3 & 5 & -2 & 0 \\ 0 & 0 & 0 & 0 & 5 & 3 \\ 0 & 0 & 0 & 0 & 0 & -7 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

still has three pivots (circled). this matrix corresponds to a linear system in which the first two variables do not appear in any of the equations. This almost never happens in real applications!

FACT: Every matrix can be reduced to row echelon form by a sequence of row swaps & row addition.

In matrix notation, if A is $m \times n$, then we can find an $m \times m$ permutation matrix (does the row swaps) and an $m \times m$ lower-triangular matrix L (undoes the row addition needed to take $A \to U$) so that

$$PA = LU$$

where U is $m \times n$ and in row echelon form.

The proof of this fact is analogous to our development of the LU factorization using Gaussian Elimination, with some slight modifications for columns with no pivot. We illustrate the procedure next w/ an example:

Consider the linear system compatible with the augmented matrix

$$\begin{array}{ccccc|c} x & y & z & u & v & \\ \end{array}$$
$$\underbrace{\begin{bmatrix} 1 & 3 & 2 & -1 & 0 \\ 2 & 6 & 1 & 4 & 3 \\ -1 & -3 & -3 & 3 & 1 \\ 3 & 9 & 8 & 7 & 2 \end{bmatrix}}_{A} \left| \underbrace{\begin{matrix} a \\ b \\ c \\ d \end{matrix}}_{\underline{b}} \right.$$

Here, we work with a generic RHS $\underline{b} = \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix}$ to get a better understanding of solution properties. The A matrix has 4 rows & 5 columns, hence we have a linear system of 4 equations in 5 unknowns. We label them x, y, z, u, v going from col'n 1 to col'n 5

The (1,1)-entry $\neq 0$, so it can be our pivot. We eliminate the entries below it to obtain:

$$\left[\begin{array}{ccccc|c} 1 & 3 & 2 & -1 & 0 & a \\ 0 & 0 & -3 & 6 & 3 & b-2a \\ 0 & 0 & -1 & 2 & 1 & c+a \\ 0 & 0 & 2 & 4 & 2 & d-3a \end{array}\right].$$

The second column has no suitable nonzero entry for second pivot, because the 3 already lies in a row with a pivot, so we move to the 3rd col'n.

The (2,3)-entry $\neq 0$, so it becomes our second pivot. Zeroing the entries below it gives us:

$$\left[\begin{array}{ccccc|c} 1 & 3 & 2 & -1 & 0 & a \\ 0 & 0 & -3 & 6 & 3 & b-2a \\ 0 & 0 & 0 & 0 & 0 & c-\frac{1}{3}b+\frac{5}{3}a \\ 0 & 0 & 0 & 0 & 4 & d+\frac{2}{3}b-\frac{13}{3}a \end{array}\right]$$

Rows 1 and 2 have their pivots in the (1,1) & (2,3) entries, respectively. No suitable pivots are available in column 4, & so the final pivot is the "4" in the fifth column. We swap rows 3 & 4 to get:

$$\begin{array}{c} \left[\begin{array}{ccccc|c} 1 & 3 & 2 & 1 & 0 & a \\ 0 & 0 & -3 & 6 & 3 & b-2a \\ 0 & 0 & 0 & 0 & 4 & d+\frac{2}{3}b+\frac{5}{3}a \\ 0 & 0 & 0 & 0 & 0 & c-\frac{1}{3}b+\frac{5}{3}a \end{array}\right] \end{array} \begin{array}{l} \left.\begin{array}{c} \\ \\ \\ \end{array}\right\} r=3 \\ \left.\begin{array}{c} \\ \end{array}\right\} m-r=1 \end{array}$$

Each highlighted row has a single pivot, and the last row is all zeros. Our pivots are 1, −3, and 4, & we have all zeros below the staircase.

WARNING: The entries in the last column, to the right of the vertical line, can not be used as pivots. These are the RHS of the equation, but pivots can only come from the coefficient matrix.

****

IMPORTANT: the rank of a matrix is the number of pivots.

****

The notion of rank will play a central role in our study of linear algebra, & we will see many different ways of defining and interpreting the rank of a matrix. For now, we explore connections between rank and the solution set of a linear system.

Some properties of rank:

(i) For an $m \times n$ matrix, we have $0 \le r = \operatorname{rank} A \le \min\{m, n\}$

(ii) The only $m \times n$ matrix with zero rank is the zero matrix

(iii) A square $n \times n$ matrix is nonsingular if and only if it has rank n.

(iv) A matrix always has the same number of pivots, even if we perform different row operations to reduce it to row echelon form.

Properties (i)–(iii) are straightforward consequences of the definition of rank. Property (iv) is not obvious, but we'll come back to it later in the class.

Solving linear systems in echelon form

Suppose we have reduced the system to row echelon form $[U|\underline{c}]$, corresponding to the linear system $U\underline{x} = \underline{c}$.

First, we check for inconsistent equations. For example, if the i^{th} row of U is all zeros, but the corresponding entry $c_i \neq 0$ in the RHS $\underline{c}$, this would represent $0 = c_i$. In this case, it means the original system $A\underline{x} = \underline{b}$ has no solution, also called incompatible.

Recalling our previous example, we see the last row only has a solution if:

$$\frac{5}{3}a - \frac{1}{3}b + c = 0$$

$$\left[\begin{array}{ccccc|c} 1 & 3 & 2 & 1 & 0 & a \\ 0 & 0 & -3 & 6 & 3 & b-2a \\ 0 & 0 & 0 & 0 & 4 & d+\frac{2}{3}b+\frac{5}{3}a \\ 0 & 0 & 0 & 0 & 0 & c-\frac{1}{3}b+\frac{5}{3}a \end{array}\right] \begin{array}{l} \left.\vphantom{\begin{array}{c}1\\1\\1\end{array}}\right\} r=3 \\ \}\, m-r=1 \end{array}$$

Supposing now we have no such inconsistencies, let's see how to solve compatible linear systems with one or more solutions.

In a linear system $U\underline{x} = \underline{c}$ in row echelon form, we split the variables in two seperate classes:

- The variables corresponding to columns containing a pivot are called basic variables.
- The variables corresponding to columns with no pivot are called free variables.

We solve $U\underline{x} = \underline{c}$ using Back Substitution, but here our goal is to express the basic variables as a function of the free variables.

This is best illustrated using an example.

Example: Let's use our previous example, but set $a=0, b=3, c=2, d=2$, resulting in:

$$
\begin{array}{ccccc|c}
x & y & z & u & v & \\
\end{array}
$$
$$
\left[\begin{array}{ccccc|c}
1 & 3 & 2 & 1 & 0 & 0 \\
0 & 0 & -3 & 6 & 3 & 3 \\
0 & 0 & 0 & 0 & 4 & 3 \\
0 & 0 & 0 & 0 & 0 & 0
\end{array}\right]
$$

We've used pink for the basic variables, and green for the free vars.

We now need to solve the reduced system:

$$
\begin{aligned}
x + 3y + 2z - u \quad\quad &= 0 \\
-3z + 6u + 3v &= 0 \\
4v &= 3 \\
0 &= 0
\end{aligned}
$$

for the basic variables x, y, z. A little bit of algebra later gives the general solution:

$$v = \frac{3}{4}, \quad z = -1 + 2u + v = -\frac{1}{4} + 2u, \quad x = -3y - 2z + u = \frac{1}{2} - 3y - 3u.$$

As their name indicates, the free variables can be chosen arbitrarily and will still provide a solution.

In general, for an $m \times n$ system (m equations, n unknowns) of rank r, there are r basic variables and $n-r$ free variables.

There are $m-r$ all zero rows in the row echelon form of the linear system, and the system is compatible if and only if the corresponding RHS term is zero.

In summary, our discussion allows us to conclude that

Theorem: A system $A\underline{x} = \underline{b}$ of m equations in n unknowns has either:

(i) exactly one solution,
(ii) infinitely many solutions, or
(iii) no solutions.

Can you identify conditions on the rank of A that imply these cases? that are necessary? For example, case (i) occurs for compatible systems w/ rank $A = n$.

A geometric perspective

Suppose we have equations in three unknowns, of the form

$$a_{i1}x_1 + a_{i2}x_2 + a_{i3}x_3 = b_i \qquad (P)$$

The solution set to each equation (P) defines a plane P_i in 3d-space.

Therefore, when solving $A\underline{x} = \underline{b}$, we are looking for a point $\underline{x} = (x_1, x_2, x_3)$ that lies in all of the planes P_i defined by the m-equations of the form P.

Geometrically, we are looking for a point $\underline{x}$ that lies in the intersection $P_1 \cap P_2 \cap \cdots \cap P_m$. The picture below illustrates the three possible cases for 3 equations in 3 unknowns:

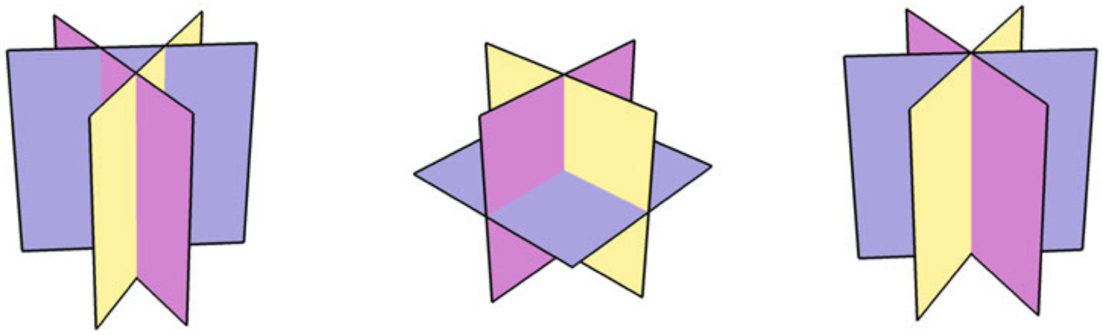

No Solution　　Unique Solution　　Infinitely Many Solutions

This ends our initial exploration of solutions of linear systems. Our story so far has been very algebraic in nature. However the power, and beauty, of linear algebra is to turn algebraic questions (does $A\underline{x} = \underline{b}$ have a soln?) into geometric ones (do the planes defined in (P) intersect).

We got a small taste of this just now, but in the next two lectures we will see much deeper and powerful connections can be made.