WEEK 2 SUPPLEMENT

1. Systems with n variables and n equations

Suppose you have a linear system with n variables and n equations - you hope it has a unique solution. Put it in matrix form $Ax = b$, where A is an $n \times n$ matrix. Compute the RREF of the augmented matrix $\left(A \mid b\right)$. If you get $\left(I_n \mid (?)\right)$, then $x = (?)$ is the solution.

If you instead have a pivot in the augmented column, your system is inconsistent—it has no solutions. If you have no pivots in the augmented column, but you have columns in the left half of the augmented matrix with no pivot, your system has infinitely many solutions. We'll discuss this in the 3rd week.

2. Inverses

If you want to figure out if an $n \times n$ matrix A is invertible (but don't care about the inverse), put A into REF. If it has n pivots, it is invertible.

If you have a 2×2 matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$, it is invertible as long as $ad - bc \neq 0$. The inverse is then $\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad-bc}\begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$.

If you want to invert an $n \times n$ matrix A, compute the RREF of $\left(A \mid I_n\right)$. If all pivots are in the left half of this $n \times 2n$ matrix, then

$$\operatorname{rref}\left(A \mid I_n\right) = \left(I_n \mid A^{-1}\right).$$

If A is invertible, then $Ax = b$ always has a unique solution, namely, $x = A^{-1}b$.

3. $PA = LU$ Factorization

Suppose you have a linear system with n variables and m equations, and *you want to solve it many times* with the same A but with many different vectors b. This is the purpose of *LU factorization*:

(1) Find $PA = LU$, where P is a permutation matrix (later), L is lower-unitriangular, and U is upper-triangular.
(2) Solve $Lc = Pb$ using forward-substitution.
(3) Solve $Ux = c$ using backward-substitution.

Now we have

$$PAx = LUx = L(Ux) = Lc = Pb;$$

multiplying both sides by P^{-1} gives $Ax = b$.

You only need to do the 1st step once—for each subsequent b vector, you can use the same L and U. This is why $PA = LU$ is so useful!

Remarks:

- *Any* matrix A has a $PA = LU$ factorization, not just square matrices. If A is an $m \times n$ matrix, then L is an $m \times m$ matrix, P is an $m \times m$ matrix, while U is an $m \times n$ matrix (the same size as A).

- The method for finding the LU factorization of a square matrix given in Jesse's lecture works, but it is inefficent, and only works for square matrices. The methods we will go over today work for all matrices.
- In theory you could solve $Ax = b$ by computing A^{-1}, so that $x = A^{-1}b$ is the solution. In practice, it is rarely a good idea to find A^{-1}!

3.1. A good algorithm for computing $A = LU$. I want to explain an algorithm for finding $PA = LU$, which is faster than finding elementary matrices, inverting, and multiplying them. This algorithm makes it clear that the task of finding $PA = LU$ is the same as doing Gaussian elimination and recording how you did it.

Let's find the $A = LU$ decomposition of $A = \begin{pmatrix} 2 & 2 & 2 \\ 1 & 2 & 3 \\ 1 & 3 & 6 \end{pmatrix}$.

	L	U
Start	$\begin{pmatrix} (?) & (?) & (?) \\ (?) & (?) & (?) \\ (?) & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} 2 & 2 & 2 \\ 1 & 2 & 3 \\ 1 & 3 & 6 \end{pmatrix}$
Eliminate in C_1	$\begin{pmatrix} 1 & (?) & (?) \\ 1/2 & (?) & (?) \\ 1/2 & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} 2 & 2 & 2 \\ 0 & 1 & 2 \\ 0 & 2 & 5 \end{pmatrix}$
Eliminate in C_2	$\begin{pmatrix} 1 & 0 & (?) \\ 1/2 & 1 & (?) \\ 1/2 & 2 & (?) \end{pmatrix}$	$\begin{pmatrix} 2 & 2 & 2 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}$.
Eliminate in C_3	$\begin{pmatrix} 1 & 0 & 0 \\ 1/2 & 1 & 0 \\ 1/2 & 2 & 1 \end{pmatrix}$	$\begin{pmatrix} 2 & 2 & 2 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}$.

Hence we have

$$L = \begin{pmatrix} 1 & 0 & 0 \\ 1/2 & 1 & 0 \\ 1/2 & 2 & 1 \end{pmatrix} \qquad U = \begin{pmatrix} 2 & 2 & 2 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}.$$

We do our row operations on the U column, and keep track of what we have done in the L column. At the end, we are left the L and U.

(1) The first step eliminates the entries below the first pivot. This uses two row operations, $R_2 -= \frac{1}{2}R_1$ and $R_3 -= \frac{1}{2}R_1$. We put a 1 in the first entry of L, and record the 1/2's below it.
(2) The next step eliminates the entries below the second pivot. This uses one row operation $R_3 -= 2R_2$. We put a 1 in the L, and record the 2 below it.
(3) The final step is just book-keeping, since there are no entries below the final pivot. Fill in the final column of L with zeros, with the last entry being 1.

Summary: Do only *subtraction* operations, and record the multipliers in the L matrix. Note that the first column of L is just the first column of U in the previous row, divided by the pivot; likewise with the second column, etc.

This method should emphasize that L keeps track of the steps you took to find U, the REF.

3.2. A good algorithm for computing $PA = LU$. In general, you will have to do row swaps in order to perform Gaussian elimination. In fact, it is usually preferable to do so, as explained in the next section. These row swaps will be recorded in a *permutation matrix* P; the resulting factorization has the form $PA = LU$.

Definition 3.3. A *permutation matrix* is a square matrix whose entries are all 0 and 1, and such that each row and each column has exactly one 1.

Fact. Any permutation matrix can be obtained by performing row swaps on the identity matrix.

Now, let's do a $PA = LU$ example. We will find P, L, and U for $A = \begin{pmatrix} 1 & 1 & 1 \\ 2 & 2 & 3 \\ 1 & 2 & 3 \end{pmatrix}$.

	P	L	U
Start	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} (?) & (?) & (?) \\ (?) & (?) & (?) \\ (?) & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ 2 & 2 & 3 \\ 1 & 2 & 3 \end{pmatrix}$
Eliminate in C_1	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & (?) & (?) \\ 2 & (?) & (?) \\ 1 & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 1 & 2 \end{pmatrix}$
Choose pivot for C_2 ($R_2 \longleftrightarrow R_3$)	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & (?) & (?) \\ 1 & (?) & (?) \\ 2 & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}$
Eliminate in C_2	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & (?) \\ 1 & 1 & (?) \\ 2 & 0 & (?) \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}$
Eliminate in C_3	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}$

$$P = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \qquad L = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix} \qquad U = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}.$$

(1) The first elimination step proceeds as before.
(2) To continue with elimination, we need to swap row 2 and row 3. Crucially, swap row 2 and row 3 in the P and L columns as well!
(3) We continue with the elimination of entries below the pivot in column 2 and 3 - there is nothing to do, and we record this in our matrix L.

Summary: When you are choosing a pivot for a column, perform the row swap on the P, L, and U matrices.

4. Maximal Partial Pivoting

This is a method of performing Gaussian elimination where we perform row swaps to make our pivot entry, at each step, as large (in absolute value) as possible. These row swaps are unnecessary for humans doing elimination by hand, but are important for a computer: without MPP, Gaussian elimination is quite numerically unstable (as we will see in the problem session). MPP is not the only possible *pivoting strategy* (way of doing extra row swaps to reduce instability), but it is the simplest.

Let's find a $PA = LU$ decomposition of this matrix using MPP:

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

	P	L	U
Start	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} (?) & (?) & (?) \\ (?) & (?) & (?) \\ (?) & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ -10 & -20 & -30 \\ 5 & 15 & 10 \end{pmatrix}$
Choose pivot for C_1 ($R_1 \longleftrightarrow R_2$)	$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} (?) & (?) & (?) \\ (?) & (?) & (?) \\ (?) & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} -10 & -20 & -30 \\ 1 & 1 & 1 \\ 5 & 15 & 10 \end{pmatrix}$
Eliminate in C_1	$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & (?) & (?) \\ -0.1 & (?) & (?) \\ -0.5 & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} -10 & -20 & -30 \\ 0 & -1 & -2 \\ 0 & 5 & -5 \end{pmatrix}$
Choose pivot for C_2 ($R_2 \longleftrightarrow R_3$)	$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & (?) & (?) \\ -0.5 & (?) & (?) \\ -0.1 & (?) & (?) \end{pmatrix}$	$\begin{pmatrix} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & -1 & -2 \end{pmatrix}$
Eliminate in C_2	$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & (?) \\ -0.5 & 1 & (?) \\ -0.1 & -0.2 & (?) \end{pmatrix}$	$\begin{pmatrix} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & 0 & -3 \end{pmatrix}$
Eliminate in C_3	$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ -0.5 & 1 & 0 \\ -0.1 & -0.2 & 1 \end{pmatrix}$	$\begin{pmatrix} -10 & -20 & -30 \\ 0 & 5 & -5 \\ 0 & 0 & -3 \end{pmatrix}$

(1) When we choose the pivot entry for C_1, we see that $|-10| > |5|, |1|$, so -10 should be our pivot entry. Thus we swap $R_1 \longleftrightarrow R_2$ to move -10 to the first row.

(2) When we choose the pivot entry for C_2, we see that $|5| > |-1|$, and so we swap $R_2 \longleftrightarrow R_3$ to move 5 to the second row.