Math 201 lecture for Monday, Week 7

Matrix inversion

Last time, we defined matrix multiplication: if A is an $m\times p$ matrix and B is a $p\times n$ matrix, then AB is the $m\times n$ matrix with i,j-entry

$$(AB)_{ij} := \sum_{k=1}^{p} A_{ik}B_{kj}.$$

If $m = n$, then BA would also be defined, but it is usually that case that $AB \neq BA$. Another peculiar thing is that for matrices, there are "zero divisors", i.e., matrices A, B such that $AB = 0$, but neither A nor B is a zero matrix. For example,

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

Diagonal matrices. The matrix A is a *diagonal matrix* if its only nonzero entries appear along the diagonal: $A_{ij} = 0$ if $i \neq j$. This terminology makes sense regardless of the dimensions of A, but is usually used in the case of square matrices, i.e., for the case where A is an $n\times n$ matrix. In that case, we write

$$A = \operatorname{diag}(a_1, \dots, a_n)$$

where $A_{ii} = a_i$ for $i = 1\dots,n$ (and $A_{ij} = 0$, otherwise.). For instance,

$$\operatorname{diag}(1,4,0,6) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{pmatrix}.$$

Identity matrices. The $n\times n$ *identity matrix* is the $n\times n$ matrix

$$I_n = \operatorname{diag}(1,\dots,1).$$

It has the following property: $AI_n = A$ and $I_nB = B$ whenever these products make sense. For instance,

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

and

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

Inverses. Let A be an $m \times n$ matrix, and let B be an $n \times m$ matrix. If $AB = I_n$, we say A is a *left-inverse* for B and B is a *right-inverse* for A. For example,

$$A = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & -1 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}.$$

Then

$$AB = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ 0 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Hence, A is a left-inverse for B and B is a right-inverse for A. On the other hand,

$$BA = \begin{pmatrix} 1 & -1 \\ 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 1 \end{pmatrix}.$$

So B is not a left-inverse of A and A is not a right-inverse of B. (In fact, B does not have a left-inverse and A does not have a right-inverse. This has to do with their ranks not being high enough. The connection with solving systems of equations we describe below explains that.)

We will mainly be interested in inverses for square matrices. Suppose that A is an $n \times n$ matrix. Suppose B is a right-inverse. So B is an $n \times n$ matrix such that $AB = I_n$. Since matrix multiplication is not commutative, the value of BA is not immediately clear. However, in fact, we have the following important result:

Theorem. Let A and B be $n \times n$ matrices. The following are equivalent:

(a) $AB = I_n$.
(b) $BA = I_n$.

If $AB = I_n$, we say A and B are *invertible* and write $A^{-1} = B$ and $B^{-1} = A$. The following are equivalent:

(i) A is invertible.
(ii) $\operatorname{rank}(A) = n$.
(iii) The reduced echelon form of A is I_n.

The proof of this theorem will follow from an elegant algorithm for computing the inverse of a matrix which we present below. The equivalence of the last to items on the list is something we already know.

Calculating the inverse. Our problem now is to determine whether an inverse for a matrix exists, and if so, to calculate that inverse. The methods we present here would also be applicable to calculating right- and left-inverses of non-square matrices—it boils down to solving systems of linear equations, after all—but we will concentrate on the case of square matrices.

Example. Let

$$A = \begin{pmatrix} 0 & 3 & -1 \\ 1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}.$$

A right-inverse for A would satisfy the following:

$$\begin{pmatrix} 0 & 3 & -1 \\ 1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

So we need to find the entries $a, b, \ldots, i$. We can break this into three problems:

$$\begin{pmatrix} 0 & 3 & -1 \\ 1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \begin{pmatrix} a \\ d \\ g \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} 0 & 3 & -1 \\ 1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \begin{pmatrix} b \\ e \\ h \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} 0 & 3 & -1 \\ 1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \begin{pmatrix} c \\ f \\ i \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

Equivalently, we need to solve three systems of linear equations:

$$\begin{aligned} 0x + 3y - z &= 1 \\ x + 0y + z &= 0 \\ x - y + 0z &= 0 \end{aligned}$$

$$\begin{aligned} 0x + 3y - z &= 0 \\ x + 0y + z &= 1 \\ x - y + 0z &= 0 \end{aligned}$$

$$\begin{aligned} 0x + 3y - z &= 0 \\ x + 0y + z &= 0 \\ x - y + 0z &= 1 \end{aligned}$$

Their augmented matrices would like:

$$\left(\begin{array}{ccc|c} 0 & 3 & -1 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & -1 & 0 & 0 \end{array}\right), \quad \left(\begin{array}{ccc|c} 0 & 3 & -1 & 0 \\ 1 & 0 & 1 & 1 \\ 1 & -1 & 0 & 0 \end{array}\right), \quad \left(\begin{array}{ccc|c} 0 & 3 & -1 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & -1 & 0 & 1 \end{array}\right).$$

The row operations needed to determine the solvability of this system are the same in all three cases. So we can combine all three of these systems at once in one "super"-augmented matrix calculation:

$$\left(\begin{array}{ccc|ccc} 0 & 3 & -1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & -1 & 0 & 0 & 0 & 1 \end{array}\right) \xrightarrow{r_1 \leftrightarrow r_2} \left(\begin{array}{ccc|ccc} 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 3 & -1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 & 1 \end{array}\right)$$

$$\xrightarrow{r_3 \to r_3 - r_1} \left(\begin{array}{ccc|ccc} 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 3 & -1 & 1 & 0 & 0 \\ 0 & -1 & -1 & 0 & -1 & 1 \end{array}\right)$$

$$\xrightarrow[r_3 \to -r_3]{r_2 \leftrightarrow r_3} \left(\begin{array}{ccc|ccc} 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 & -1 \\ 0 & 3 & -1 & 1 & 0 & 0 \end{array}\right)$$

$$\xrightarrow{r_3 \to r_3 - 3r_2} \left(\begin{array}{ccc|ccc} 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 & -1 \\ 0 & 0 & -4 & 1 & -3 & 3 \end{array}\right)$$

$$\xrightarrow{r_3 \to -r_3/4} \left(\begin{array}{ccc|ccc} 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 & -1 \\ 0 & 0 & 1 & -1/4 & 3/4 & -3/4 \end{array}\right)$$

$$\xrightarrow[r_2 \to r_2 - r_3]{r_1 \to r_1 - r_3} \left(\begin{array}{ccc|ccc} 1 & 0 & 0 & 1/4 & 1/4 & 3/4 \\ 0 & 1 & 0 & 1/4 & 1/4 & -1/4 \\ 0 & 0 & 1 & -1/4 & 3/4 & -3/4 \end{array}\right).$$

Going back to the original systems of equations, we see that we need

$$\begin{pmatrix} a \\ d \\ g \end{pmatrix} = \begin{pmatrix} 1/4 \\ 1/4 \\ -1/4 \end{pmatrix}, \quad \begin{pmatrix} b \\ e \\ h \end{pmatrix} = \begin{pmatrix} 1/4 \\ 1/4 \\ 3/4 \end{pmatrix}, \quad \begin{pmatrix} c \\ d \\ i \end{pmatrix} = \begin{pmatrix} 3/4 \\ -1/4 \\ -3/4 \end{pmatrix}.$$

In other words, the following matrix is a right-inverse for A:

$$\begin{pmatrix} 1/4 & 1/4 & 3/4 \\ 1/4 & 1/4 & -1/4 \\ -1/4 & 3/4 & -3/4 \end{pmatrix}.$$

The argument we've just given for a particular matrix easily generalizes to give the following algorithm.

Algorithm for computing the inverse of a matrix. Let A be an $n \times n$ matrix. Perform row operations on the "super"-augmented matrix $[A \mid I_n]$ to compute the reduced echelon form of A:

$$(A \mid I_n) \longrightarrow \left(\tilde{A} \mid B\right).$$

There are two possibilities: either $\tilde{A} = I_n$ or not. If $\tilde{A} = I_n$, then $B = A^{-1}$. Next, we consider what happens when $\tilde{A} \neq I_n$. Since B is derived by performing row operations on I_n, we have $\mathrm{rank}(B) = \mathrm{rank}(I_n) = n$. Thus, B cannot have a row of zeros. If $\tilde{A} \neq I_n$, it must have a row of zeros. It follows that the system of equations is inconsistent, and A has no inverse.

Now suppose that rank(A) so that

$$(A \mid I_n) \longrightarrow (I_n \mid B) \tag{1}$$

and $AB = I_n$. Consider trying to find C so that $BC = I_n$. In this case, reverse the row operations in (1) to get

$$(B \mid I_n) \longrightarrow (I_n \mid A)\,,$$

and thus, $C = A$, i.e., $BA = I_n$.

In summary:

- If $\tilde{A} = I_n$ (equivalently, $\mathrm{rank}(A) = n$) then $AB = BA = I_n$. (So $B = A^{-1}$ and $A = B^{-1}$.)
- If $\tilde{A} \neq I_n$ (equivalently, $\mathrm{rank}(A) < n$), then $\tilde{A}$ has a row of zeros and A has no inverse.

In particular: $A \in M_{n \times n}$ is invertible if and only if $\mathrm{rank}(A) = n$.